The Anatomy of a Leak: AS9121

or

How We Learned to Start Worrying and Hate the Maximum Prefix Limits

Alin C. Popescu, Brian J. Premore, Todd Underwood

{alin,bj,todd}@renesys.com

Renesys Corporation

Christmas Eve Leak

- **24 Dec 2004**: 100K+ routes leaked from AS9121 (**TTnet**), globally propagated
- Bad routes resulted in misdirected/lost traffic for tens of thousands of networks: **serious global vulnerability**
- Best common practices were insufficient to prevent direct and collateral damage
- Will examine the timeline, assess the damage, and what steps operators may take for infrastructure integrity assurance

A Full Table of ... Turkey

- AS9121(**TTnet**) announces an (almost) full table to peers, including AS6762 (**Telecom Italia**)
- AS6762 has one misconfigured session with no *maximum prefix* set, so they accept 100K+ prefixes
- AS6762 propagates those prefixes to their peers, hitting *maximum prefix* limits on all of those sessions
- “Bad” prefixes originated by AS9121 replace those originated by the real owners

Sample Organizations with Hijacked Routes

Blue Cross Blue Shield of Iowa

Thomson Financial Services

Citicorp Global Information Network

MetLife Capital Corp

Pitney Bowes Credit Corporation

Brown Brothers Harriman & Company

LaSalle Partners

Kuwait Fund for Arab Economic Development

Two Events: Timeline #1

- 09:19:57 UTC 24 Dec 2004: AS9121 starts announcing 106K+ prefixes to peers
- 09:19:57: AS6762 starts carrying 106K+ prefixes originated by AS9121
- 09:19:58: Renesys hears reports of “bad” paths from 13 peers

Two Events: Timeline #1 (cont'd)

- 09:20:07: 1/3 of Renesys peers heard and believed "bad" paths
- 09:20:27: "Bad" paths spread across the Internet
- 09:36:10: Peak in announcement rate
- 10:03:00: First event ends, but AS9121 continues to announce bad prefixes throughout the rest of the day

Two Events: Timeline #2

- 19:47:06: AS9121 begins announcing bad prefixes at a high rate
- 19:47:39: Peak in announcement rate
- 19:50:00: Second event ends, but AS9121 continues to advertise bad prefixes for a long time

Damage Extremely Widespread – Highlights

- AS6762 carried 106606 bad prefixes
- AS1299 had *maximum prefix* to AS9121 set relatively low, but was not saved:
  - Heard only 1849 bad prefixes directly from AS9121
  - Carried a total of 10925 bad prefixes from other peers:

ASN	6762	1239	6453	701
Num Prefixes	4413	3997	2522	612

Collection Infrastructure

- Renesys operates a peering setup with
  - ≈100 peering sessions
  - peering at NOTA and LINX, multi-hop from elsewhere
  - peers on 6 continents
- “Full tables” from all peers
- Globally integrated view: rapid query of updates from all sources, not just a single collection point

Propagation of Bad Prefixes

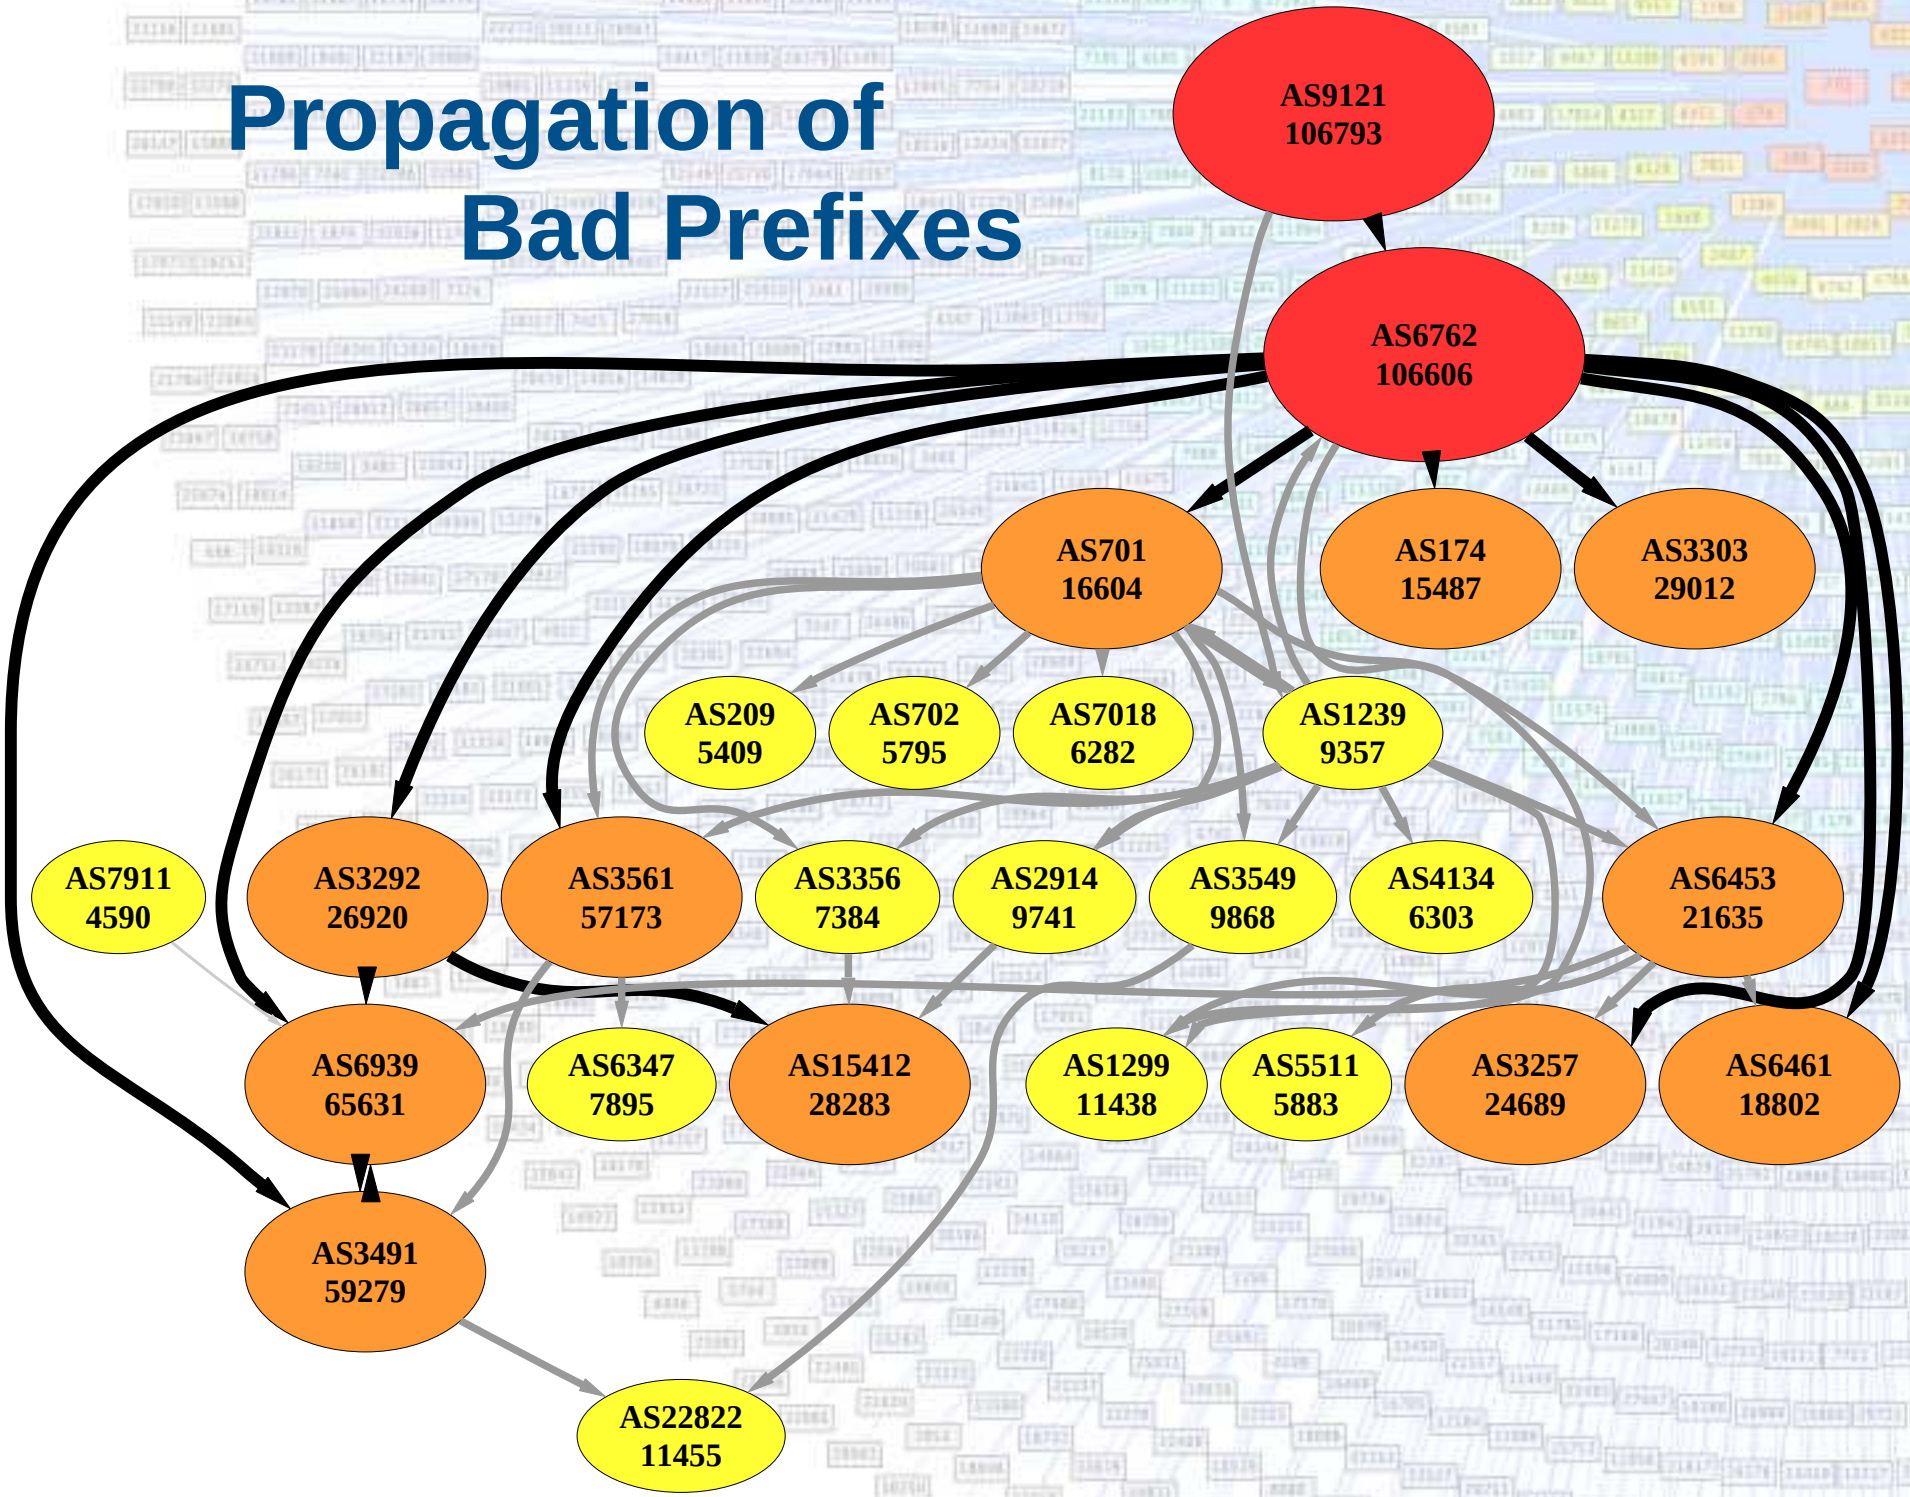

Who Spread Most Bad Prefixes

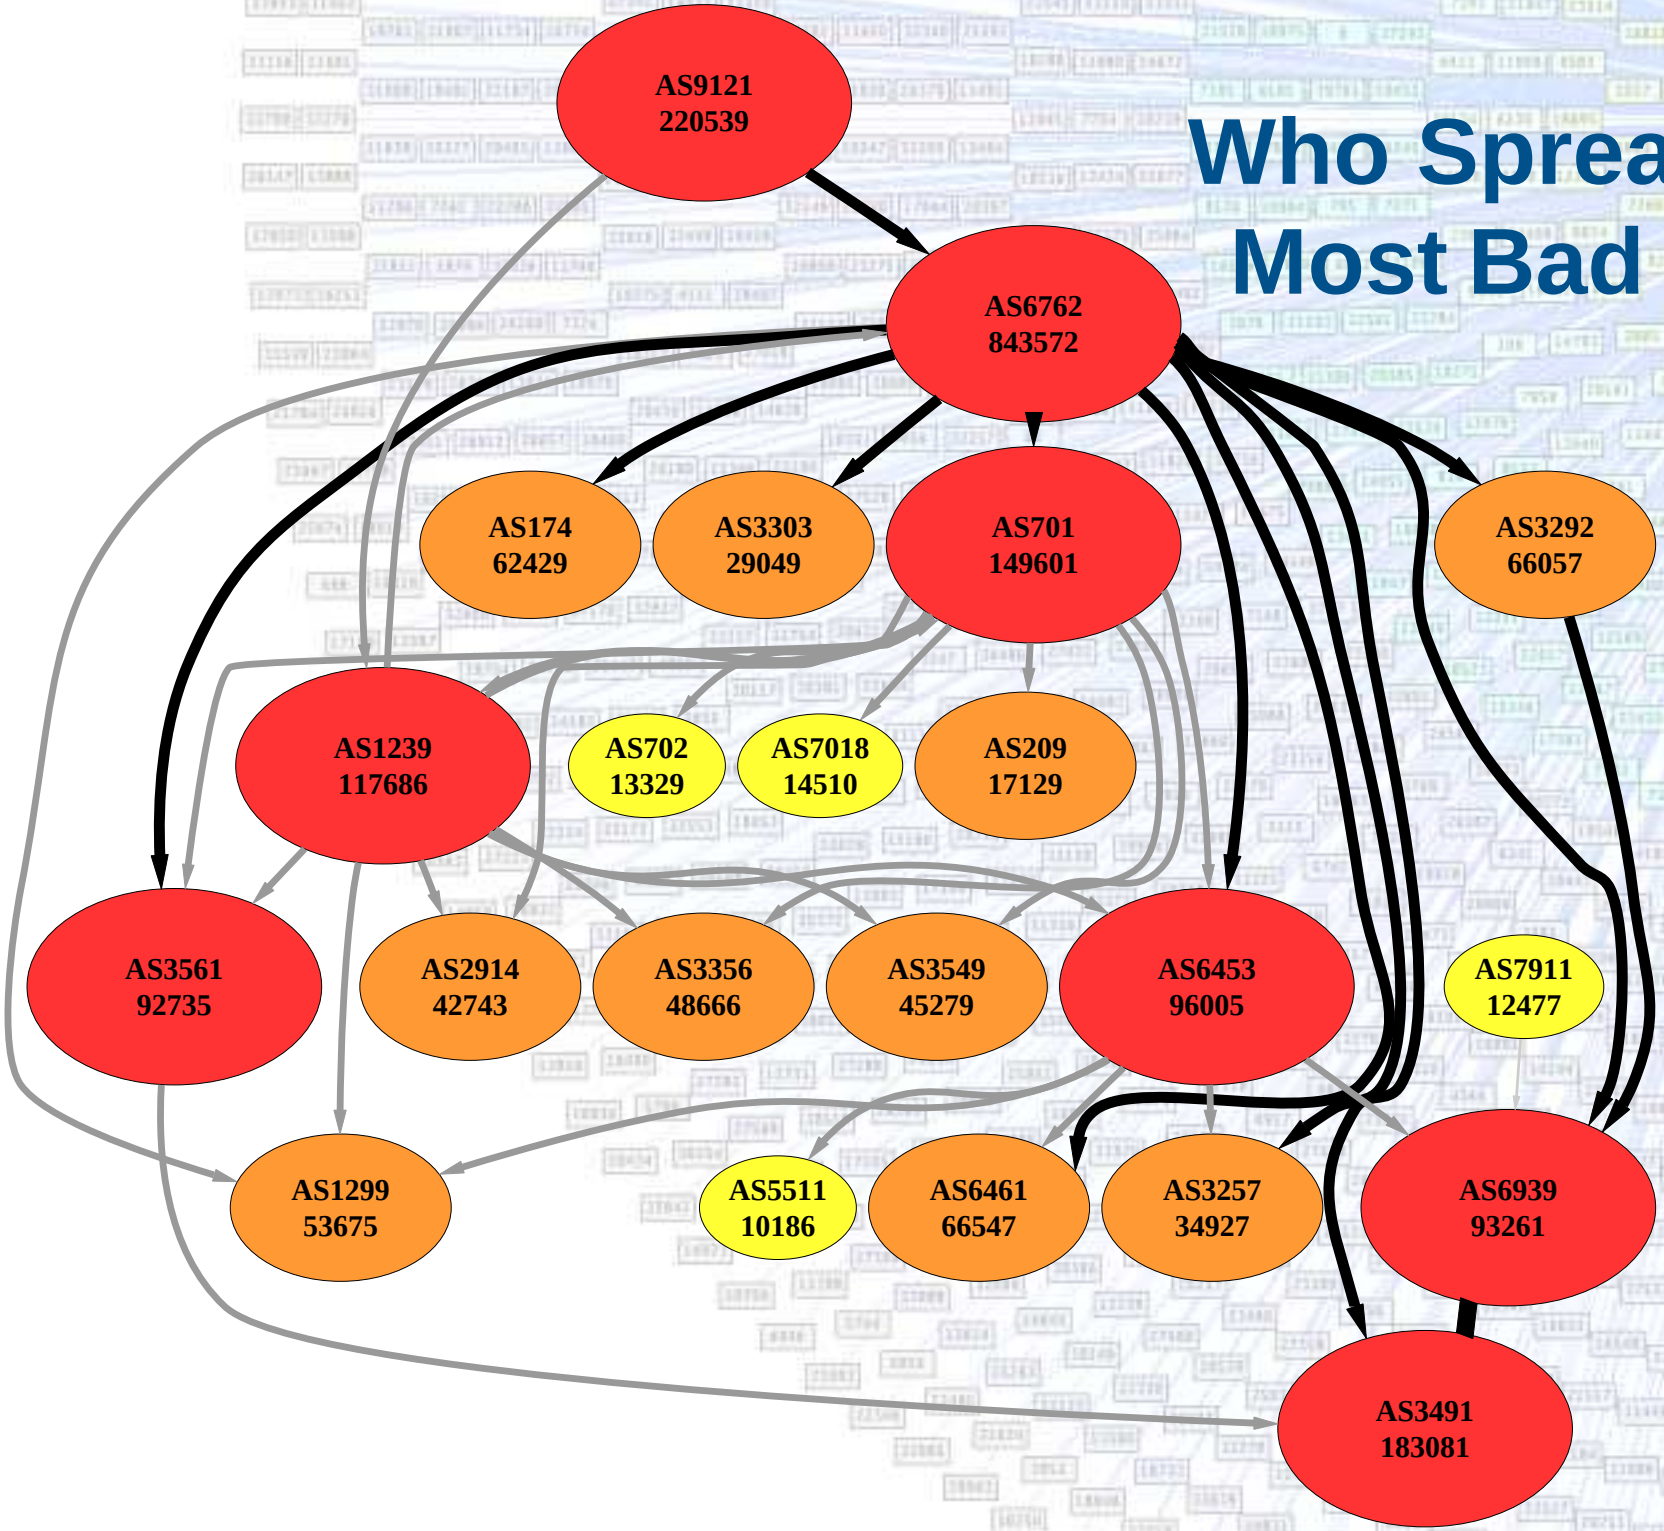

Distinct “Bad” Prefixes Over Time

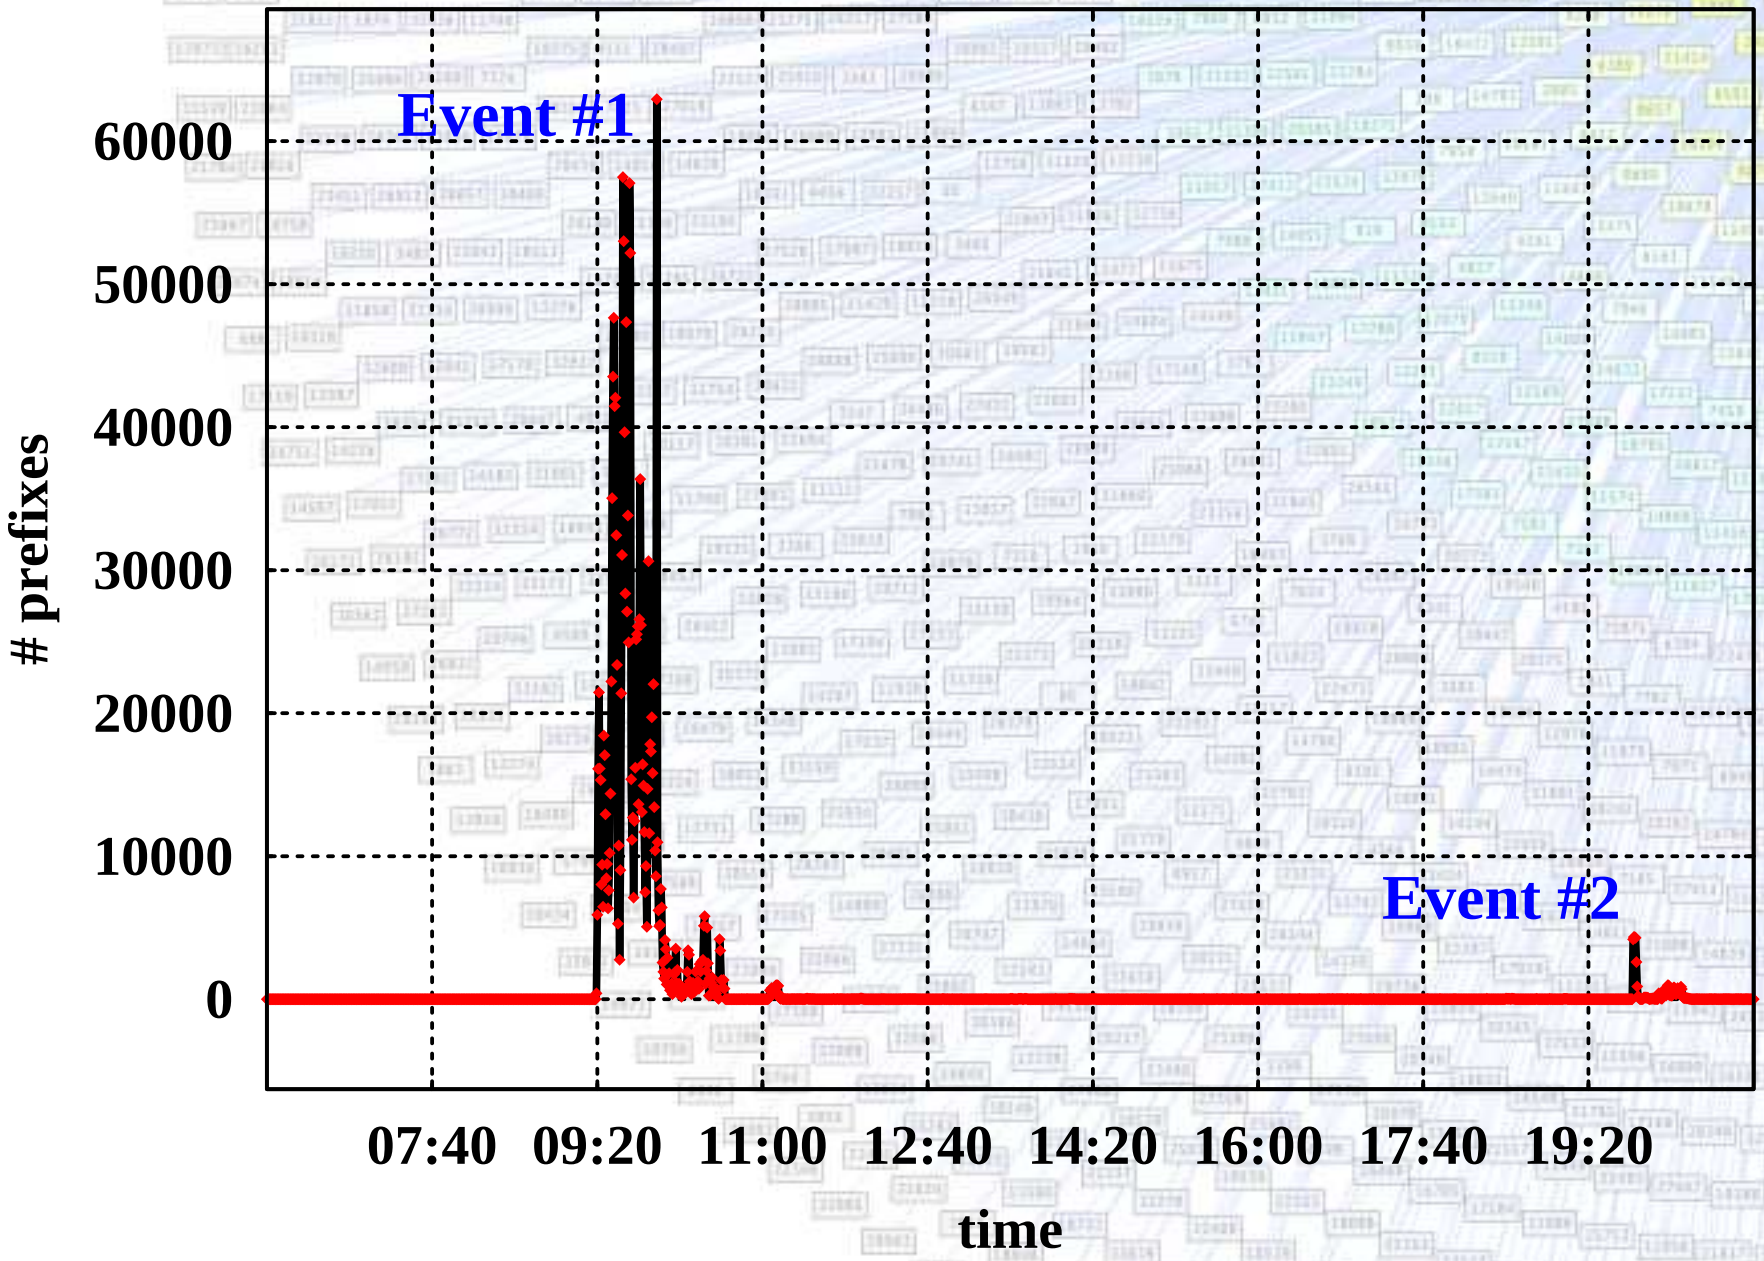

Distinct “Bad” Prefixes Over Time (log scale)

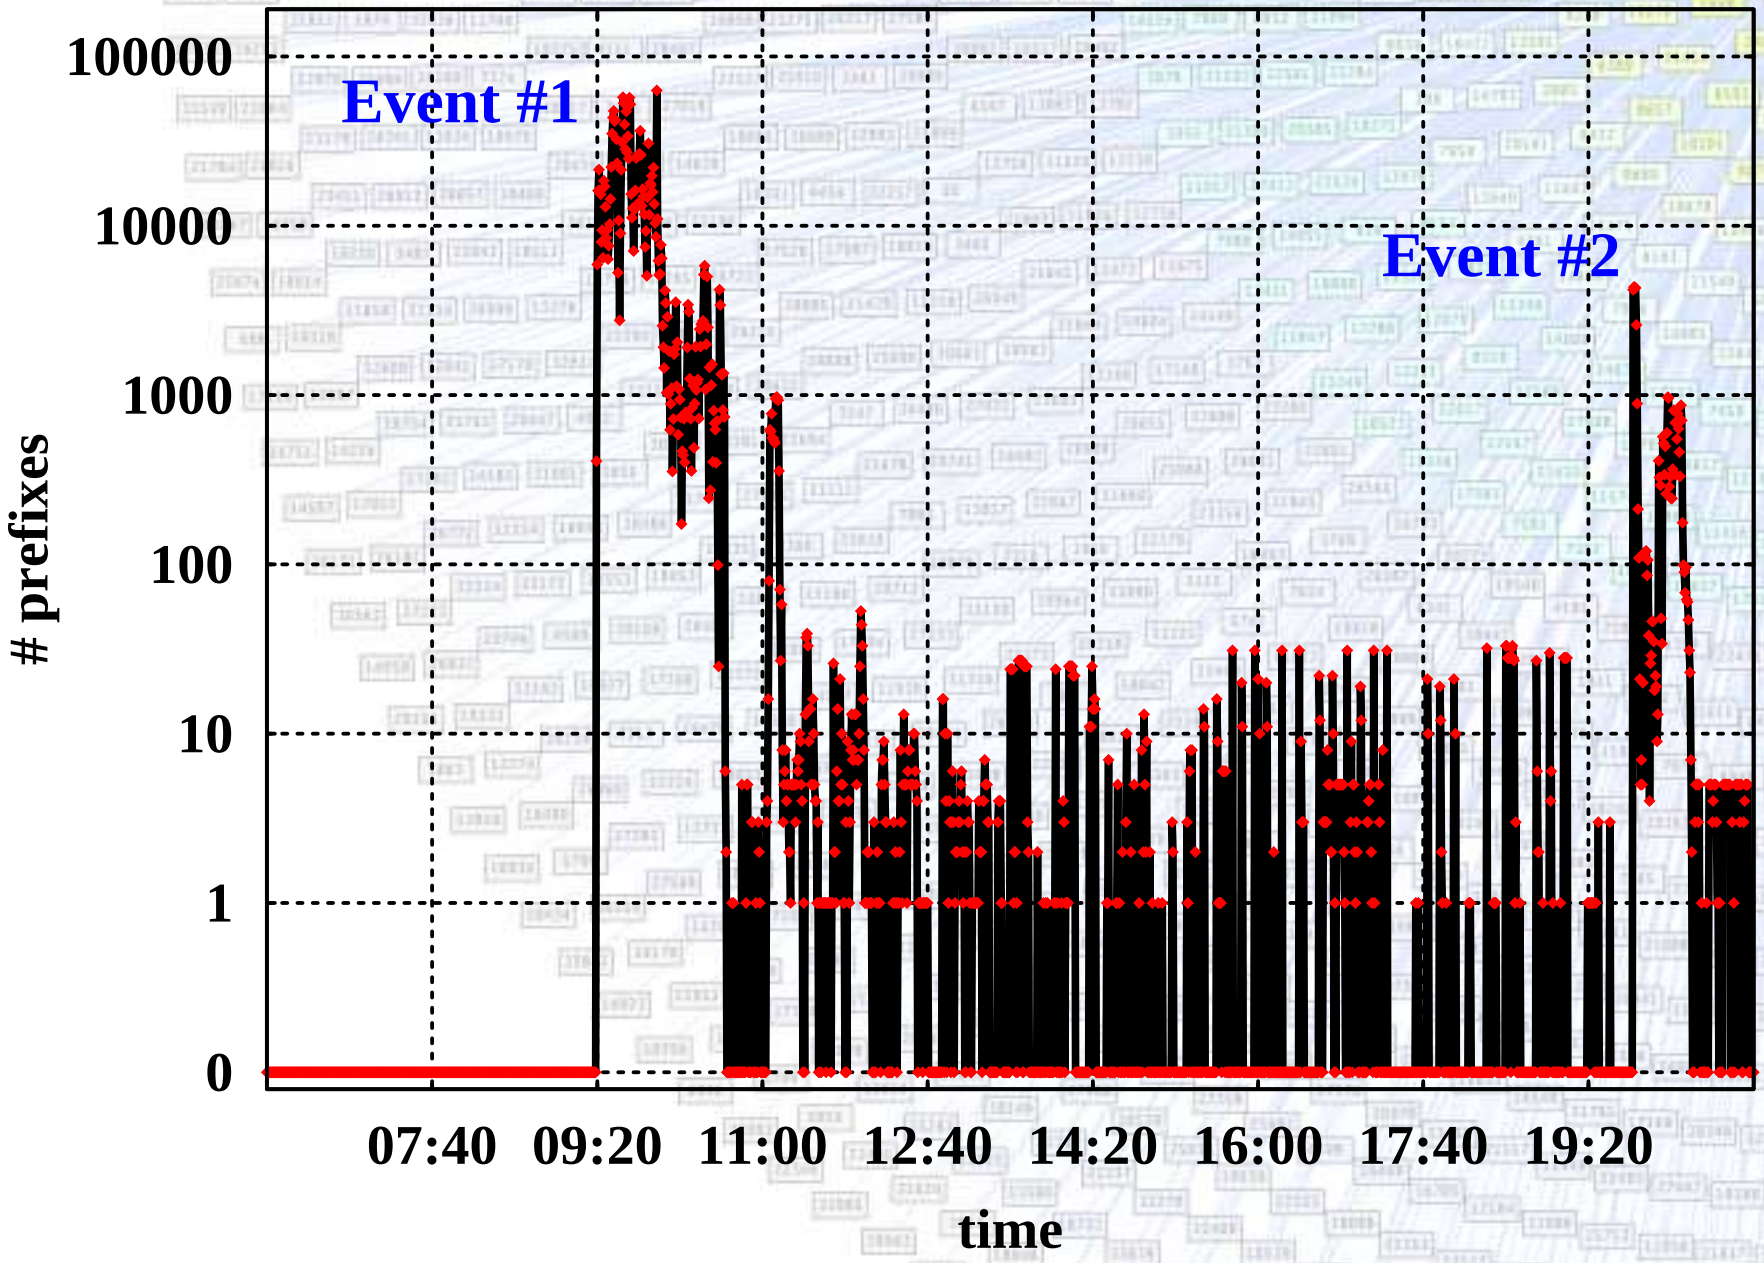

Event #1 – Zoom in

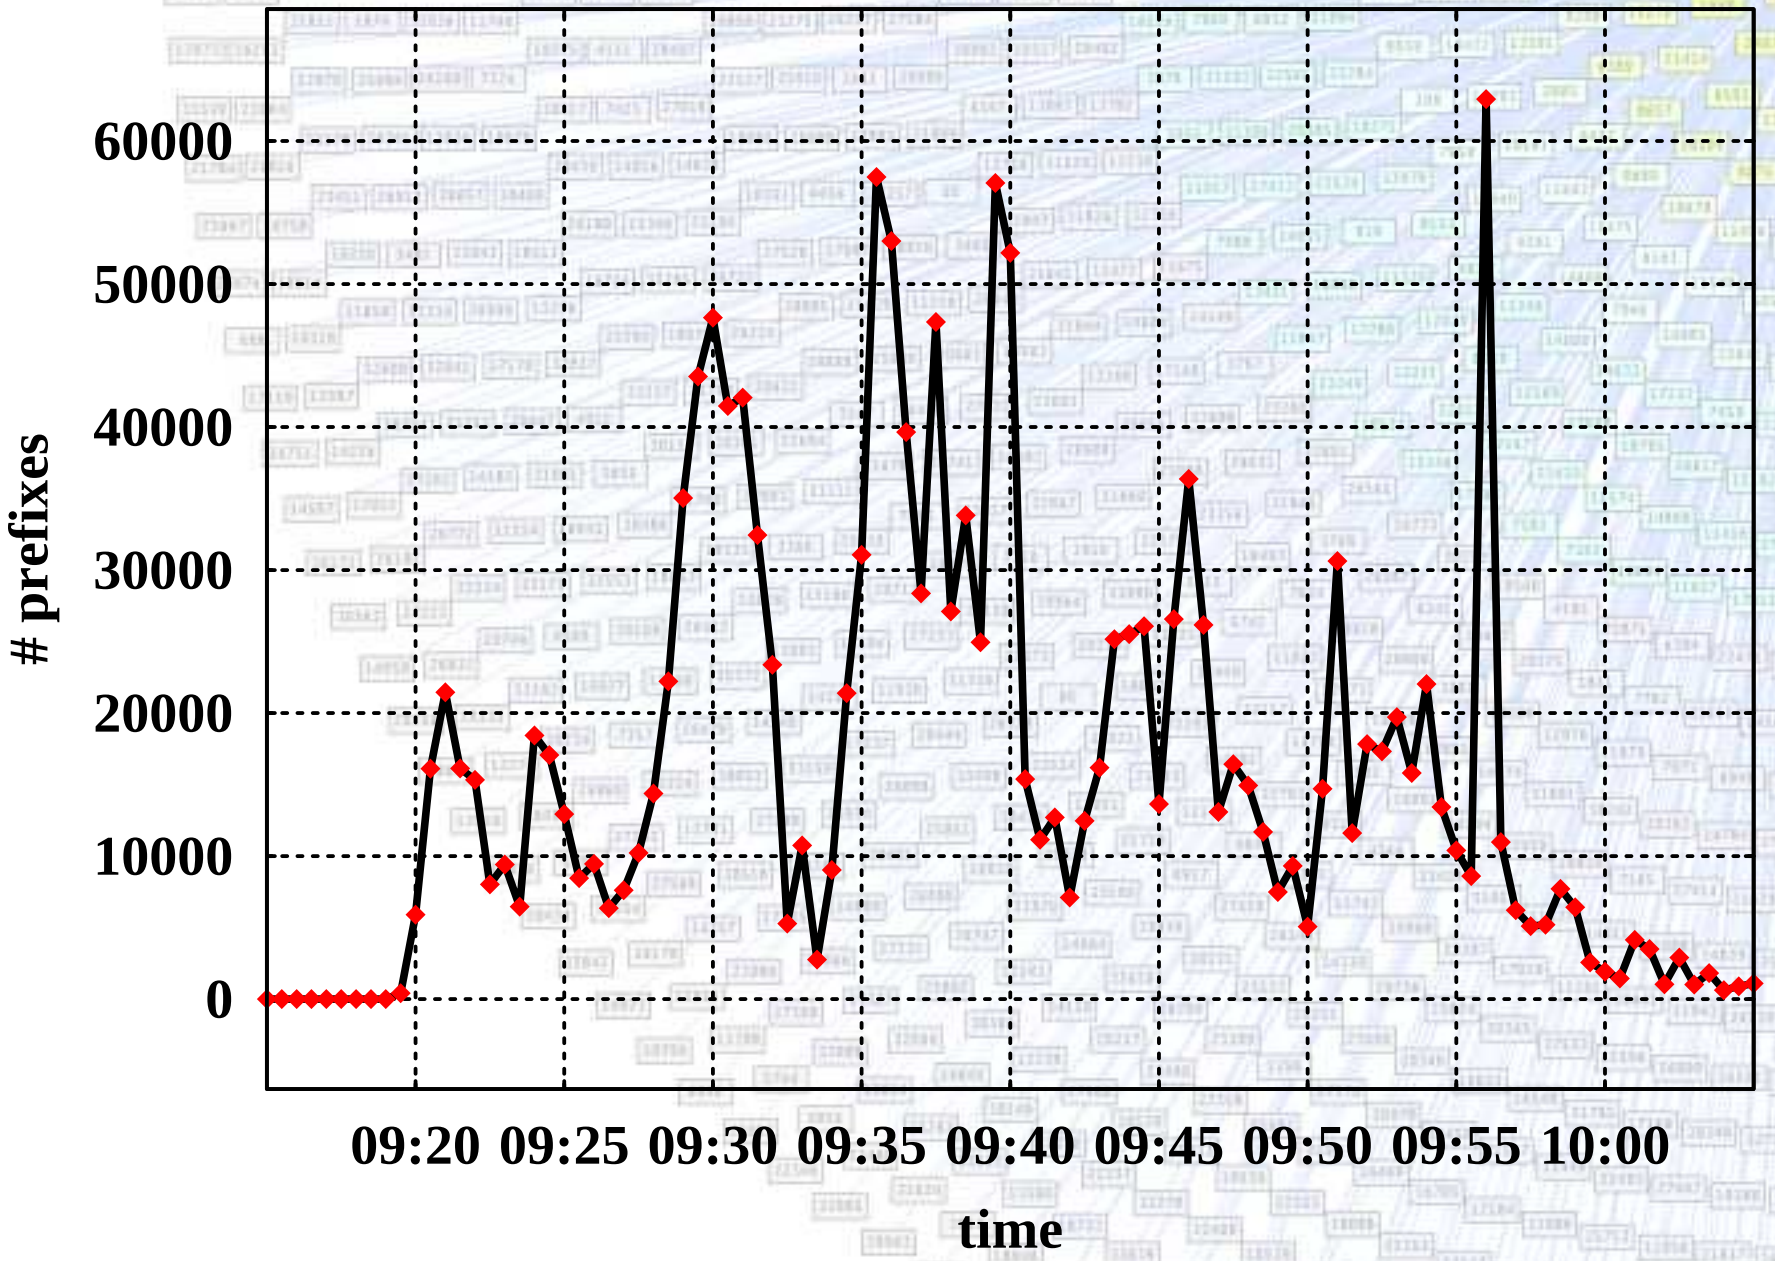

Event #2 – Zoom in

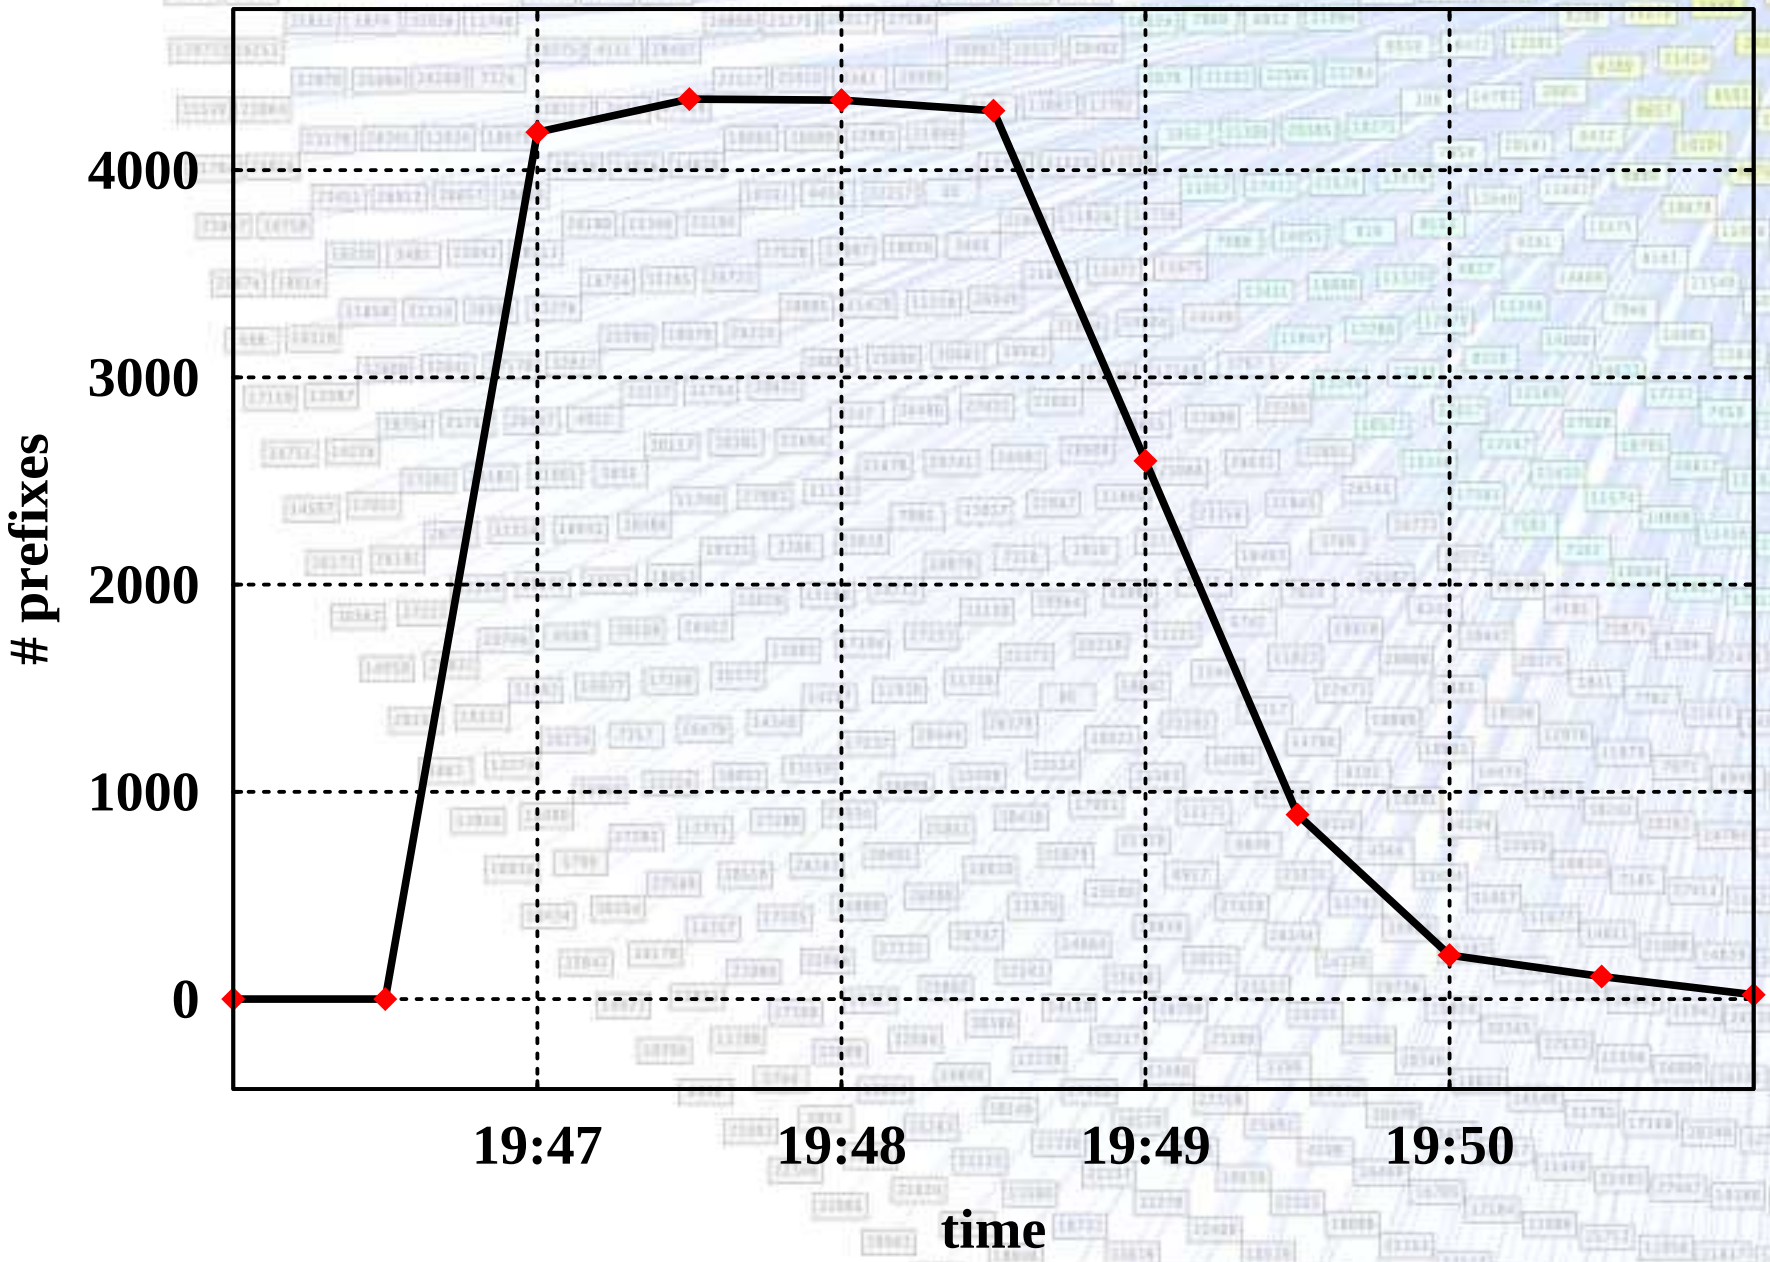

Rates of Advertisement – Event #1

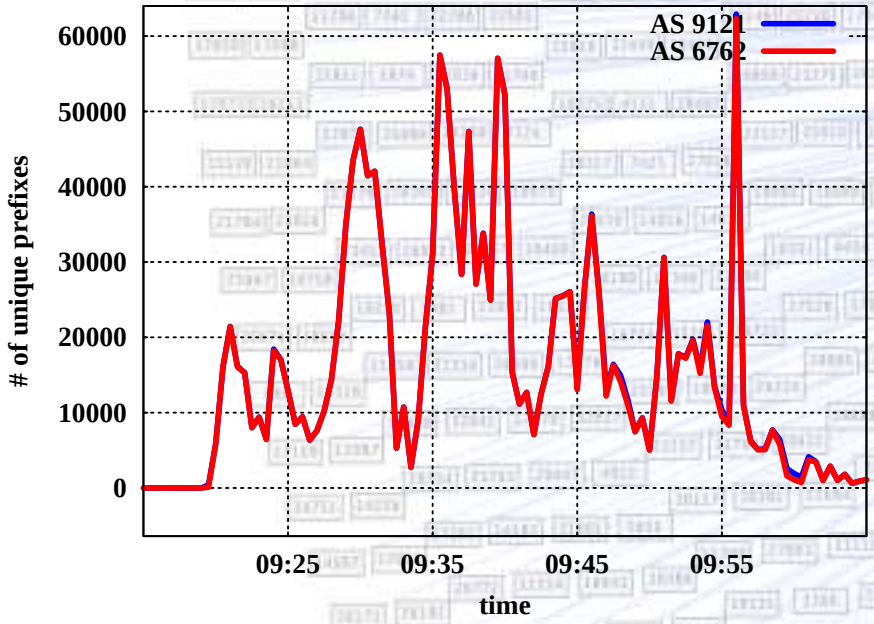

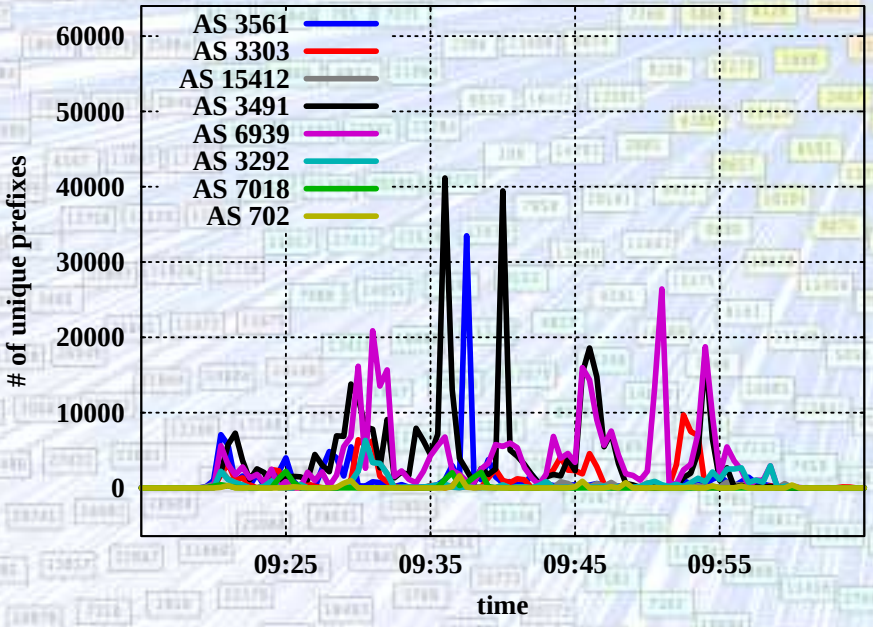

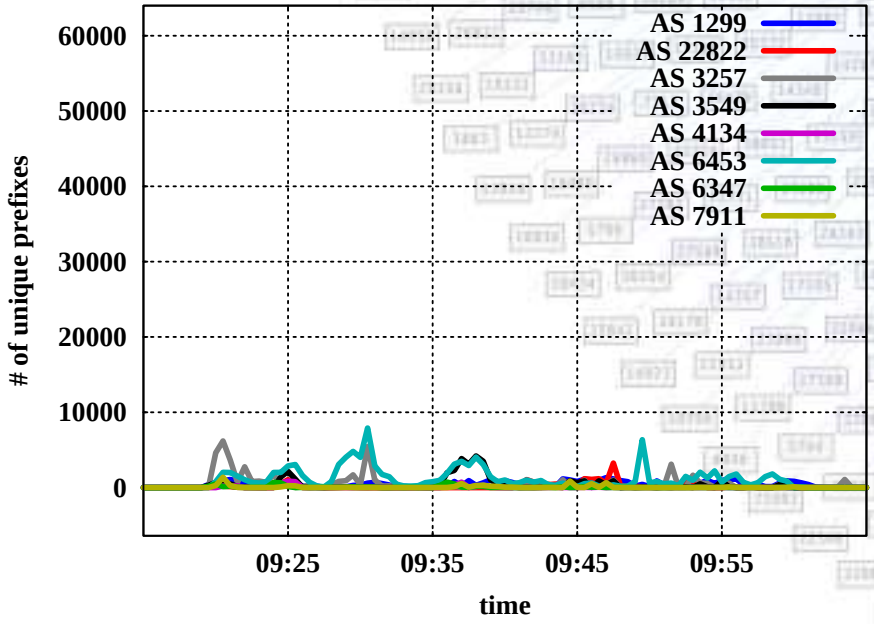

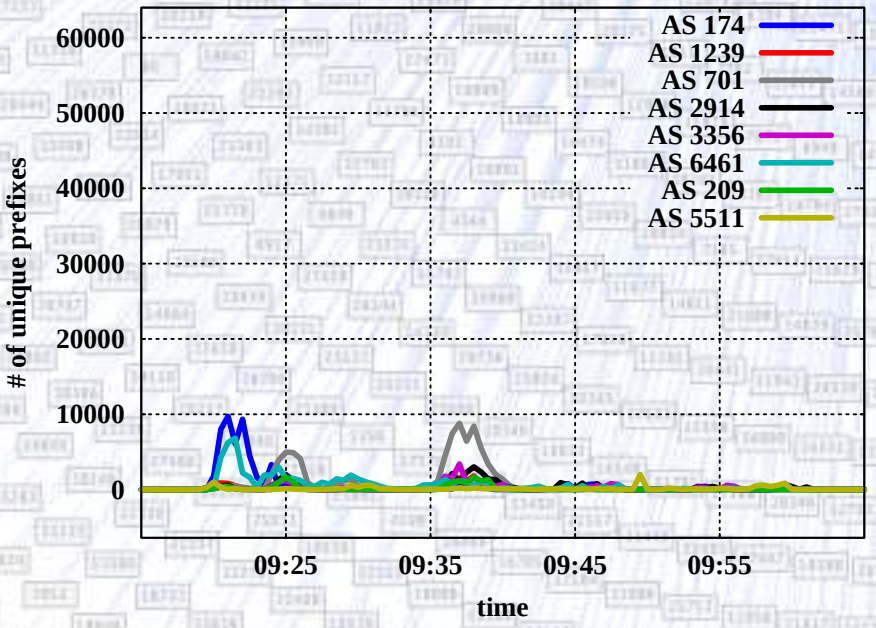

Rates of Advertisement – Event #1

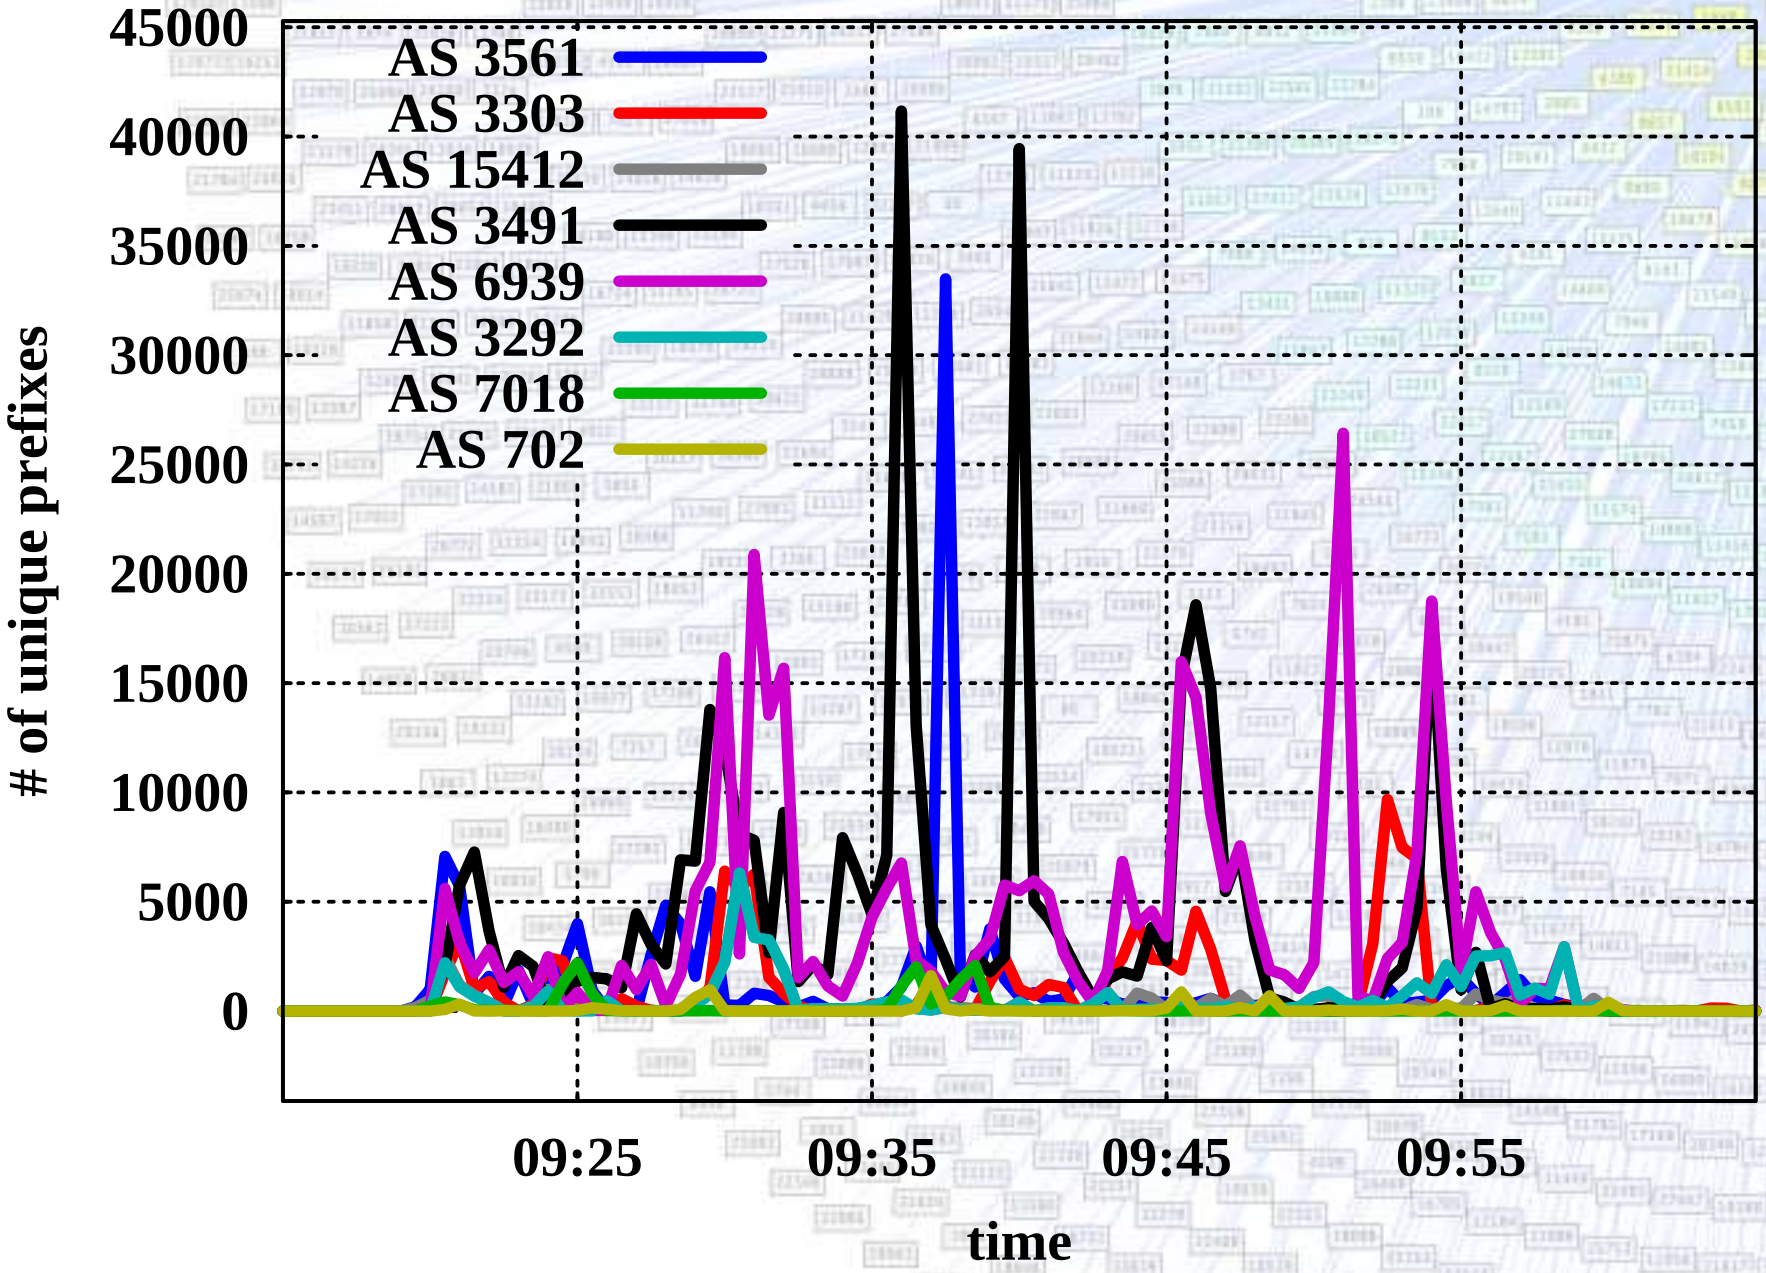

Rates of Advertisement – Event #1

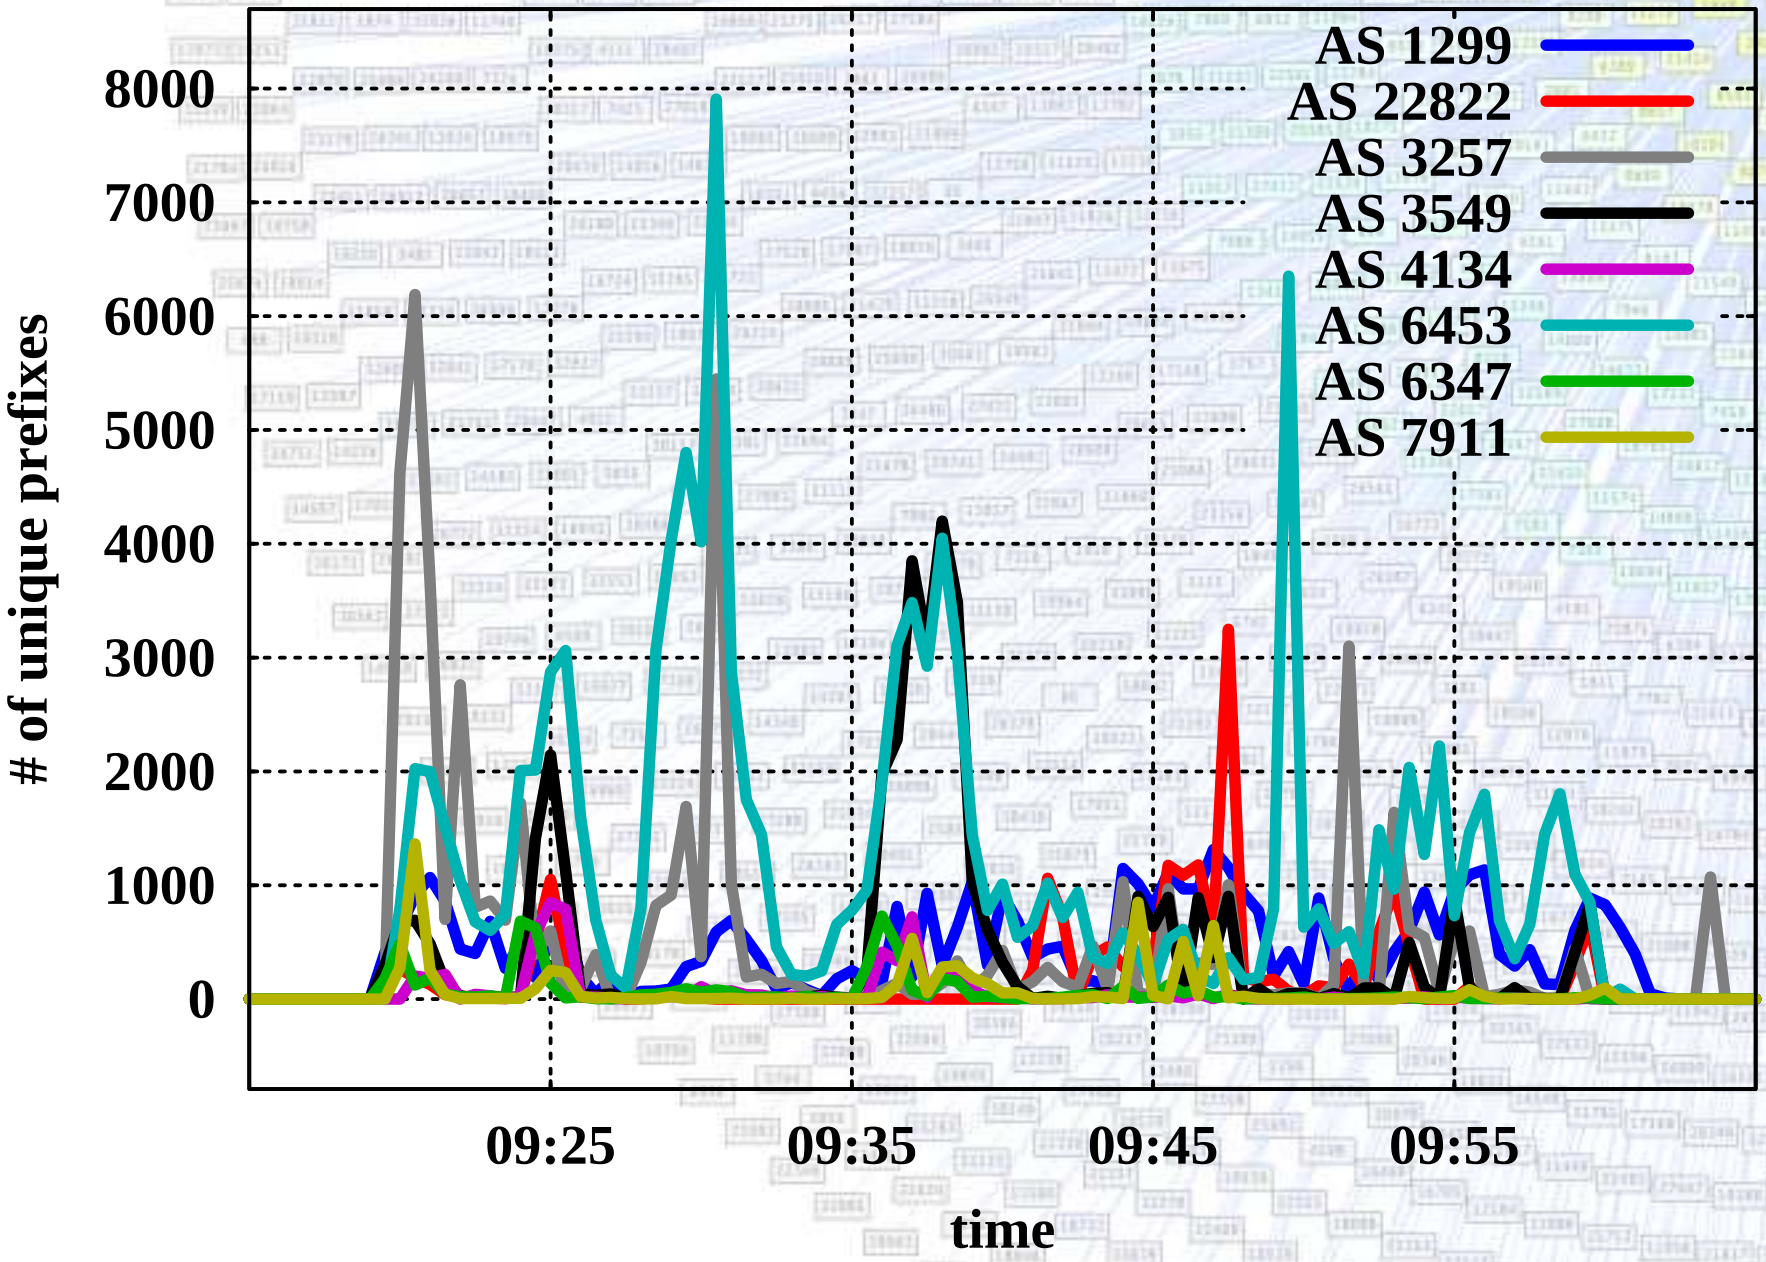

Rates of Advertisement – Event #1

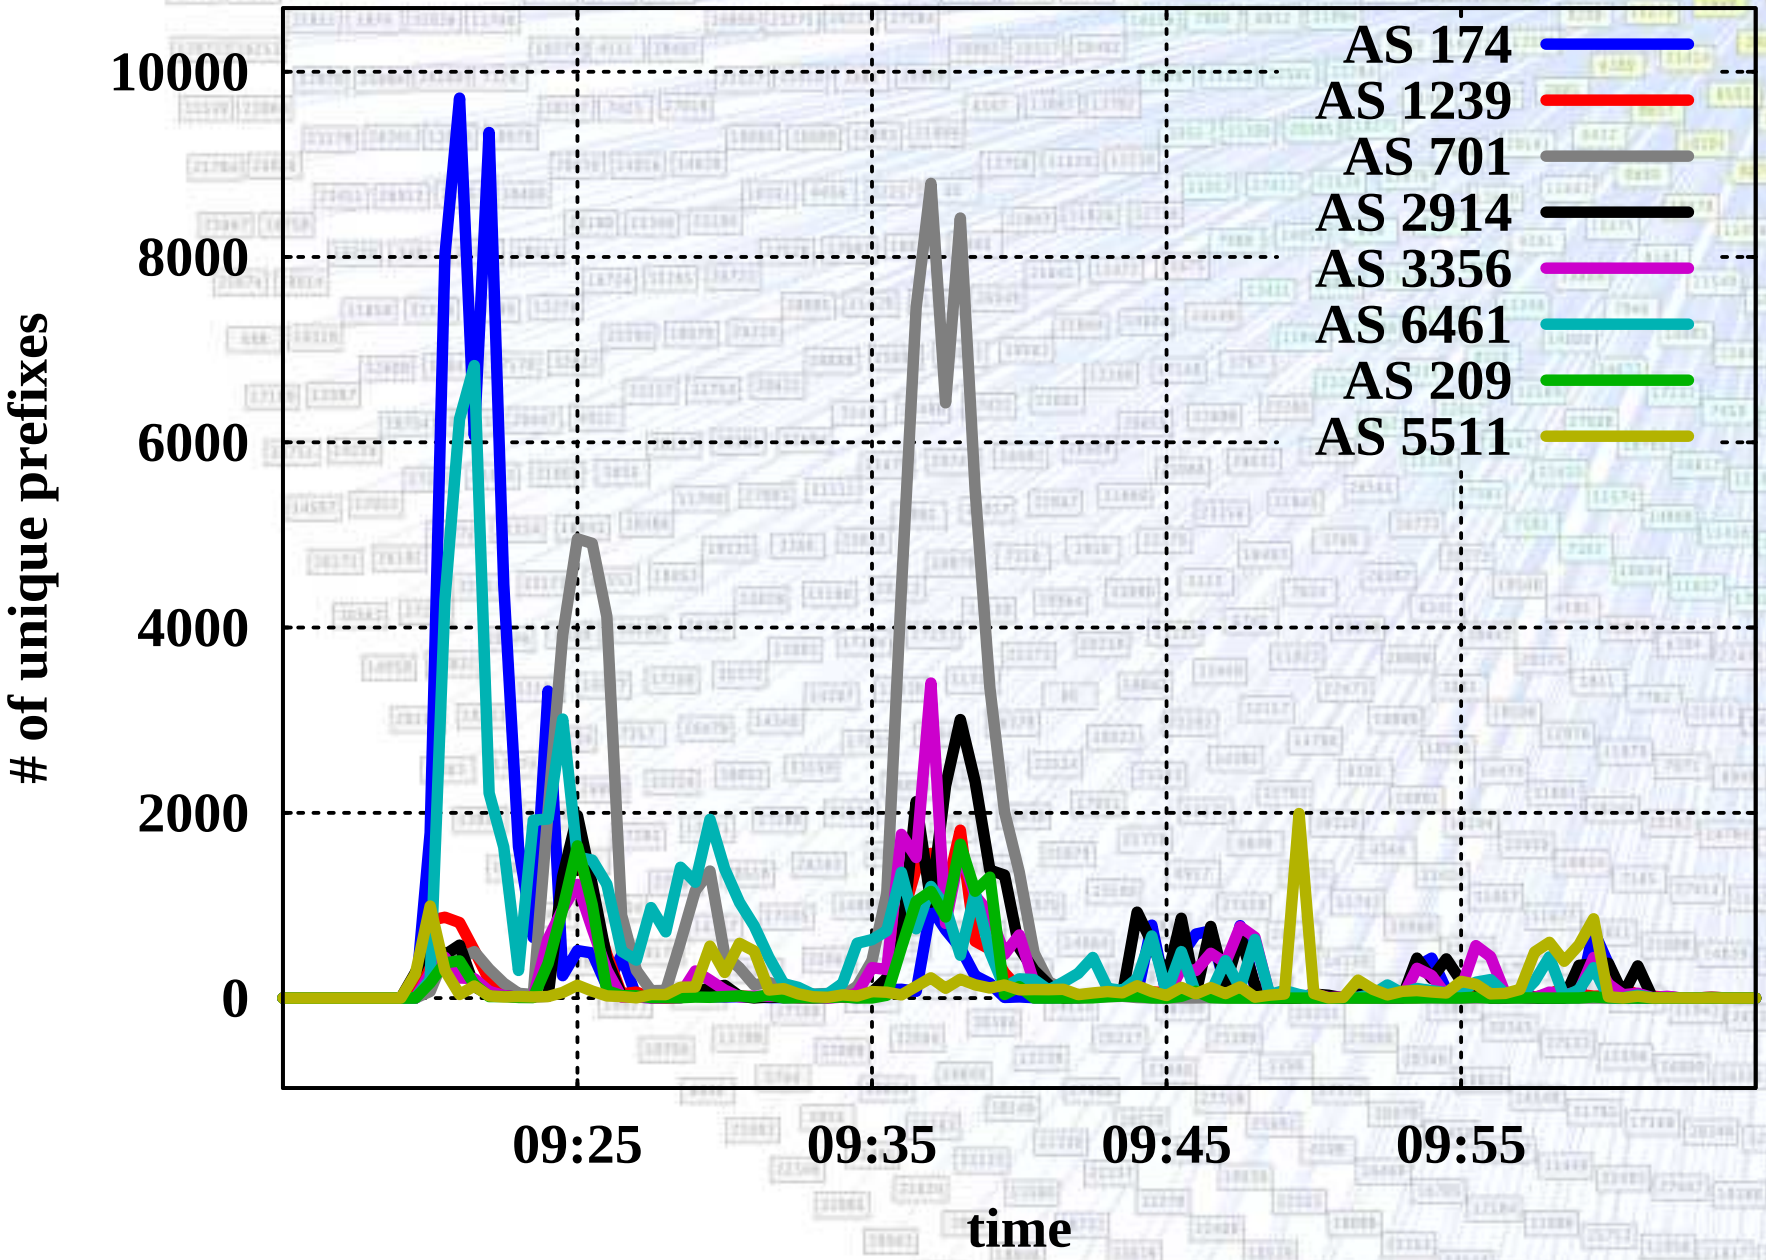

Prefixes Carried – Event #1

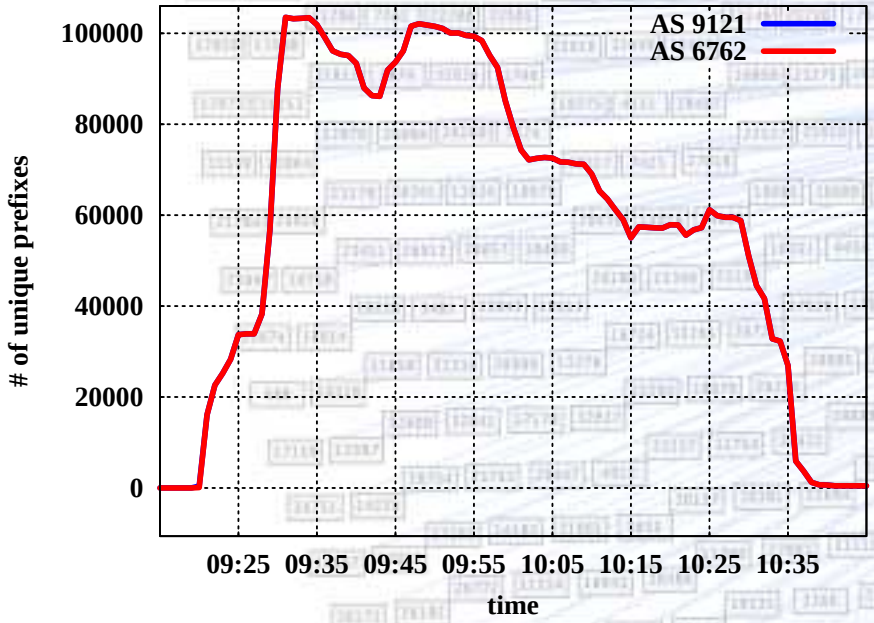

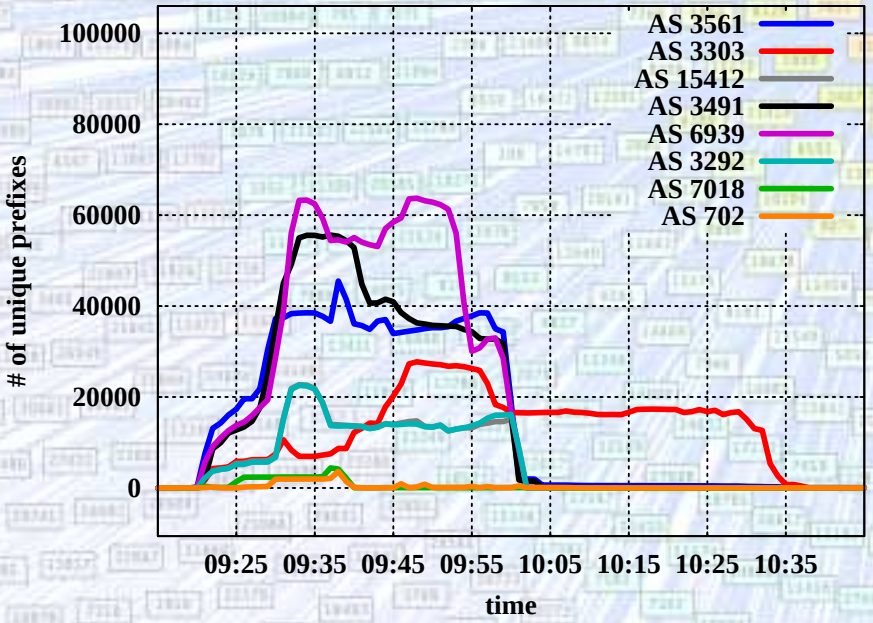

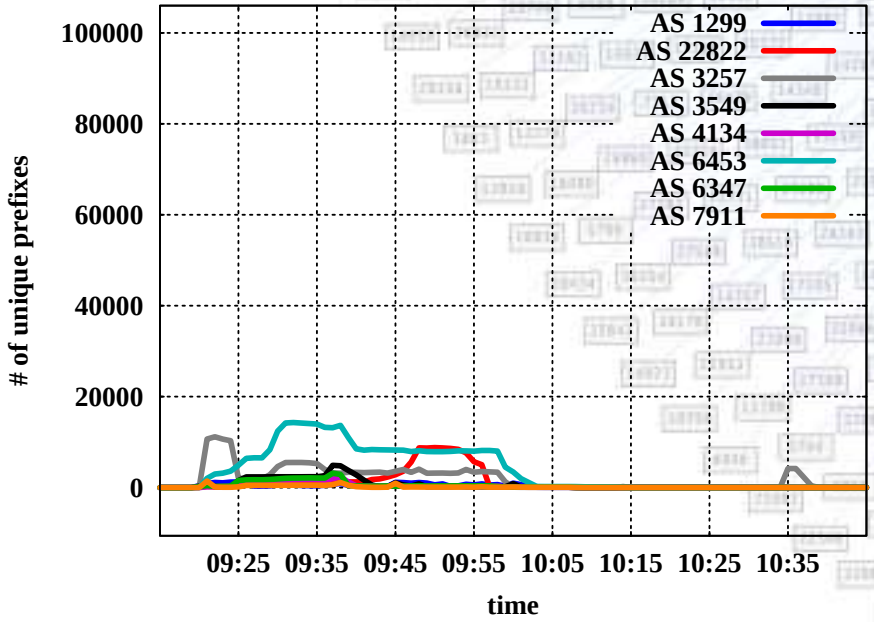

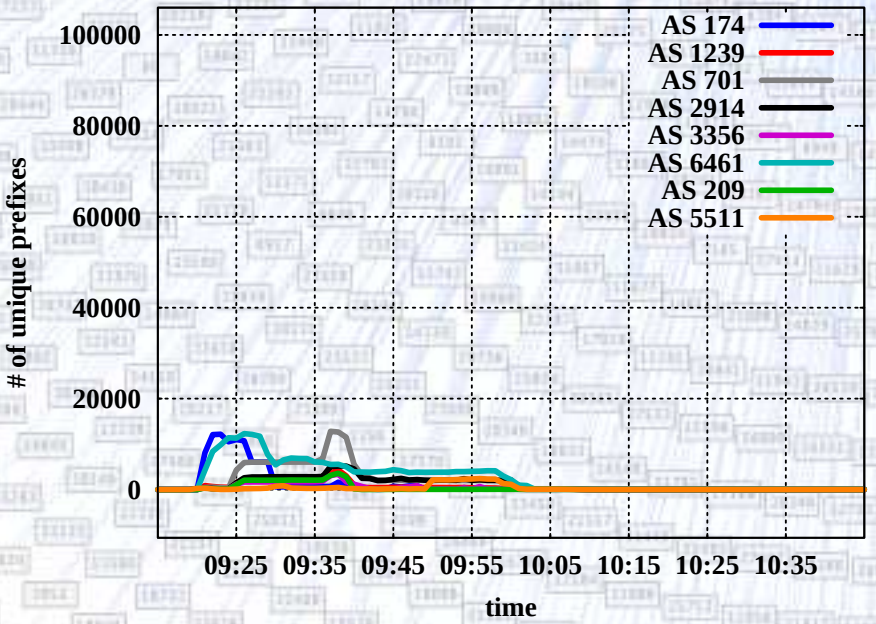

Prefixes Carried – Event #1

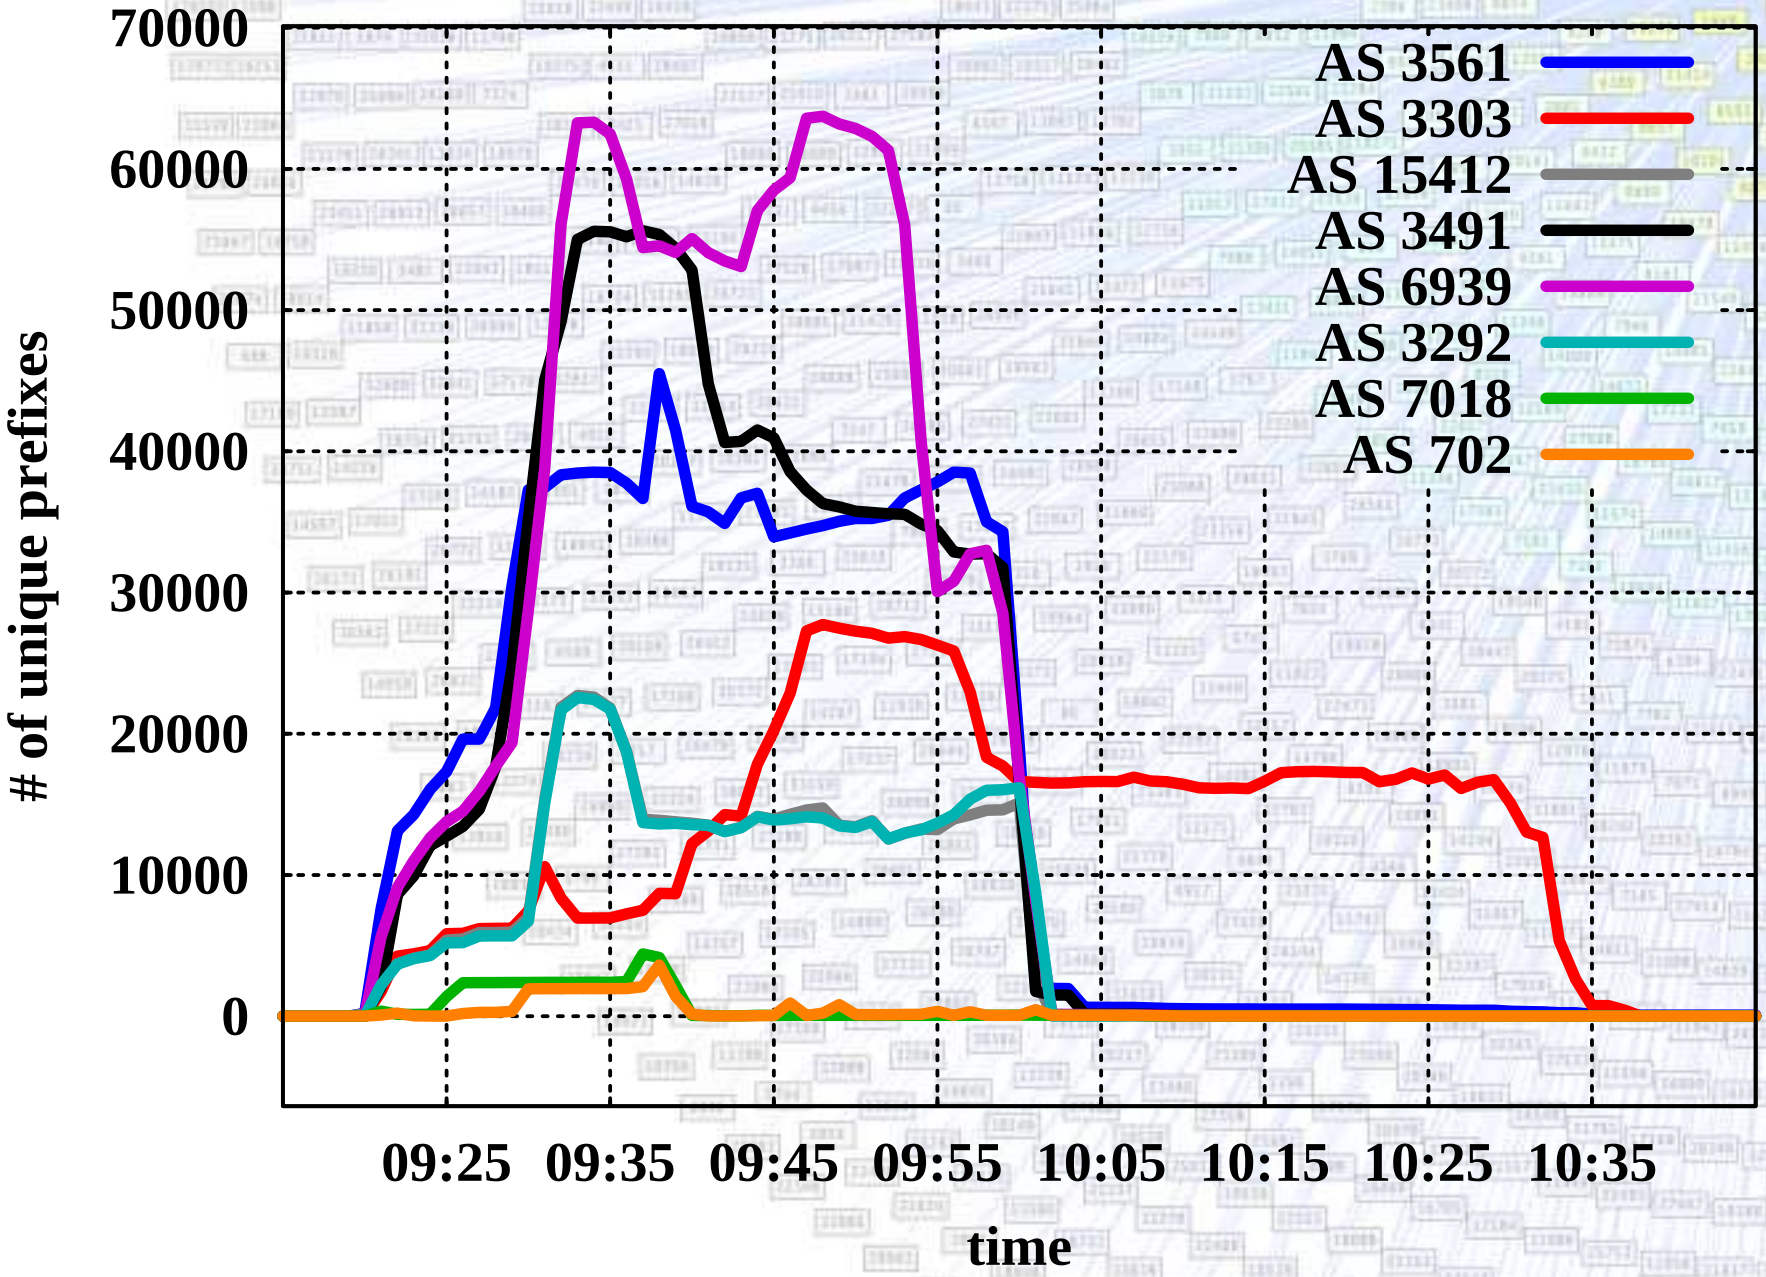

Prefixes Carried – Event #1

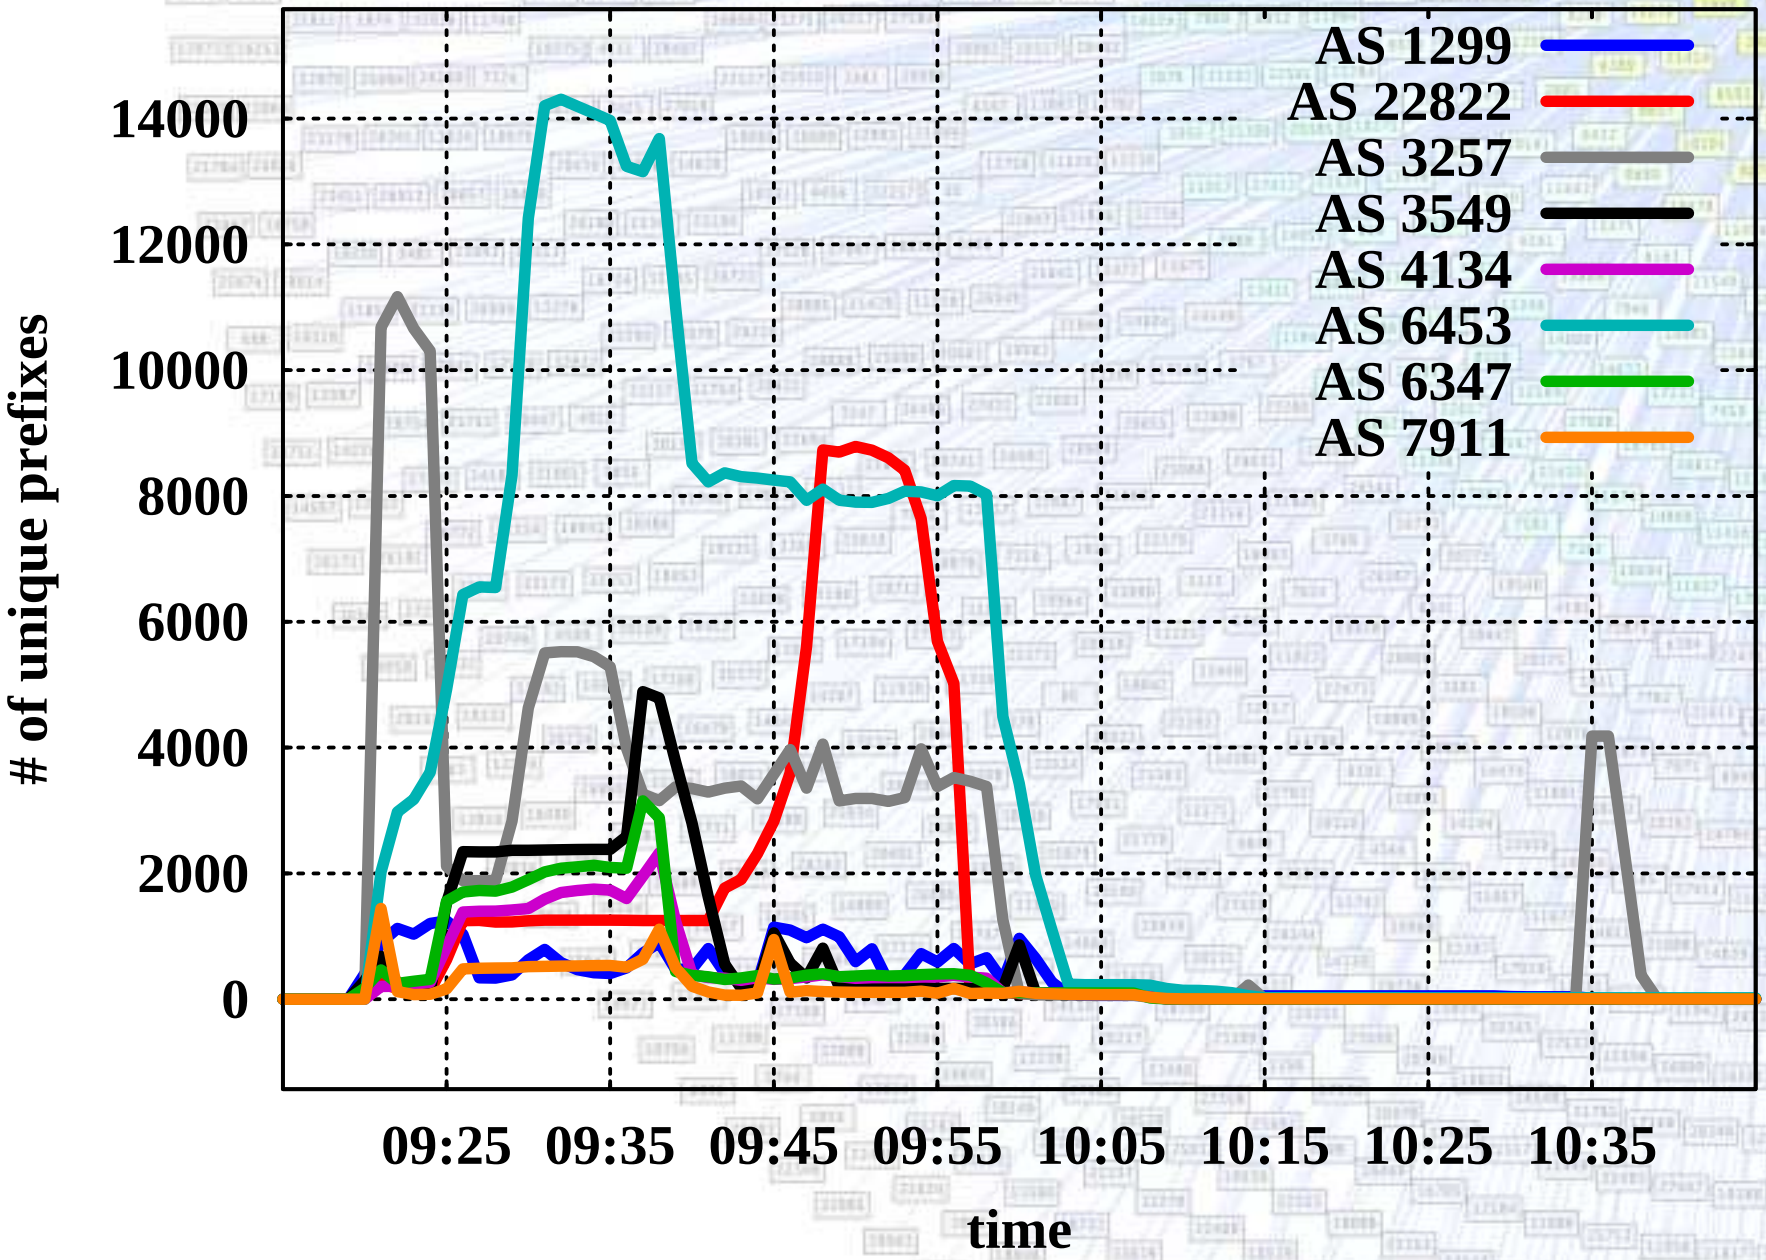

Prefixes Carried – Event #1

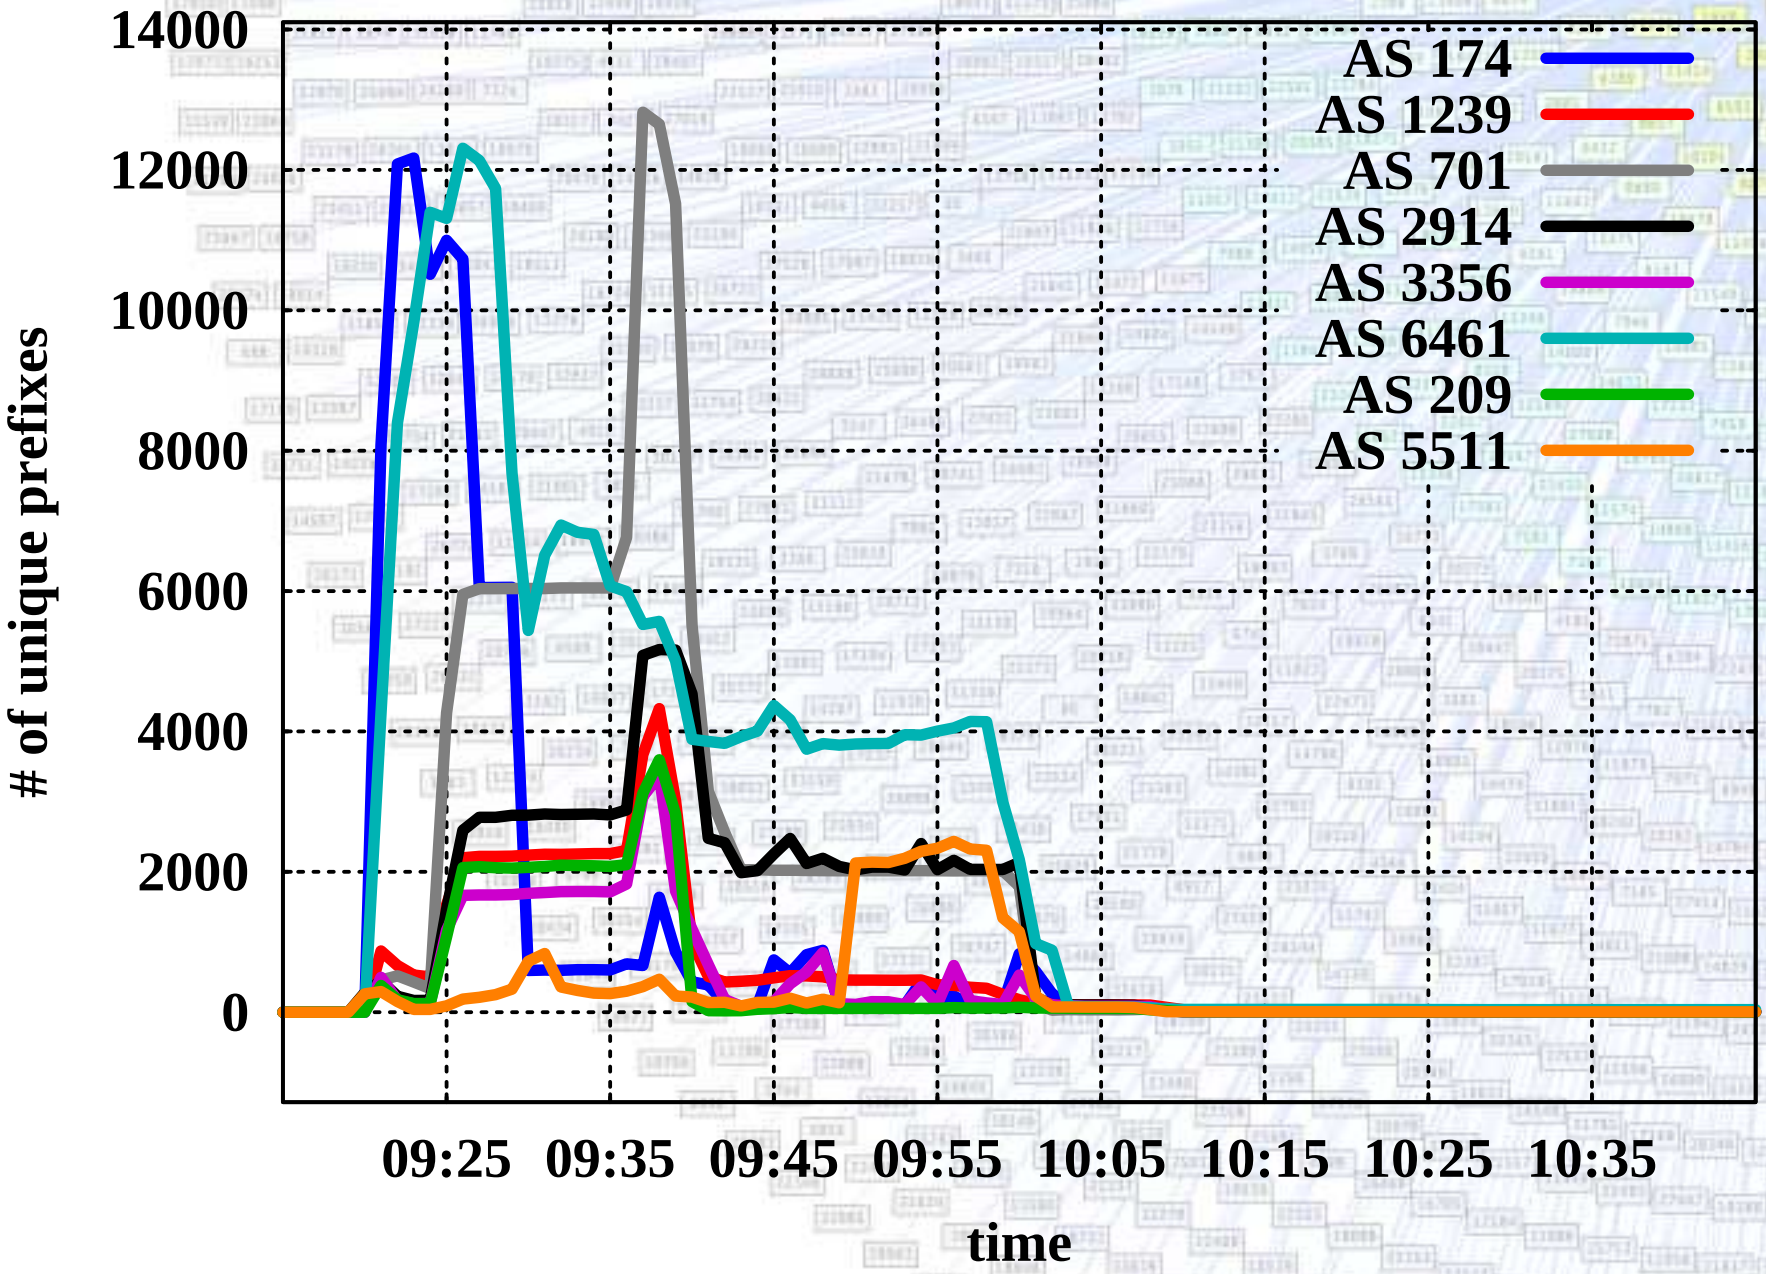

Notes on the Data

- All prefix counts are lower bounds, biased by the sampling
- It is likely that non-peer autonomous systems carried considerably more bad prefixes than what observed
- To validate the results, data from RouteViews and RIPE were also used

Operational Lessons

- Holiday staffing: not easy but matters
- Resetting a *maxpref* ’d session: should **not** be prevented by change management
- Current contact and escalation info for all peers: essential
- Tight maximum prefix settings: helps but not enough
- Transitively trusting all peers’ on-net customers: fundamentally unsafe

Future Work: Beyond maxpref

- It is impossible for large autonomous systems to prefix-filter their peers
  - Hard on some hardware: too many prefixes
  - Impossible on the people: no way to generate/maintain lists for big ASes
- It is impossible for large autonomous systems to filter on AS-path origination
  - Hard on most hardware: *regexp* 's are slow
  - Impossible on the people: no way to generate/maintain lists for big ASes
  - Wouldn't help in cases like this anyway

Future Work: Beyond maxpref

But...

- Current model is “trust all peers transitively”
- Bad things will continue to happen
- *maxpref* settings didn’t help much and won’t in the future

Therefore...

- Alternative solutions must be considered, including prefix filtering and AS-path origination filtering peers.

Thank you

peering@renesys.com

{alin,bj,todd}@renesys.com

Additional Slides

Rates of Advertisement – Event #2

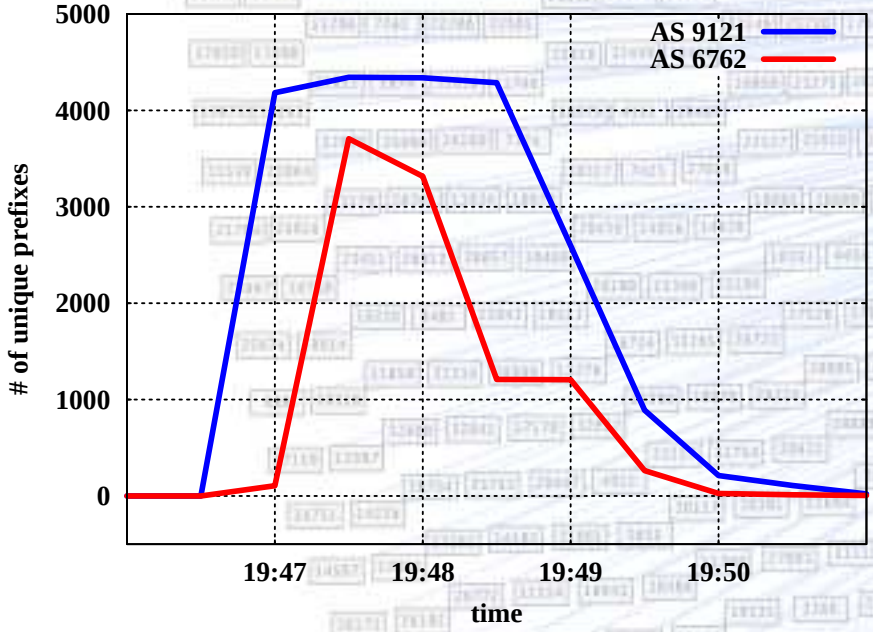

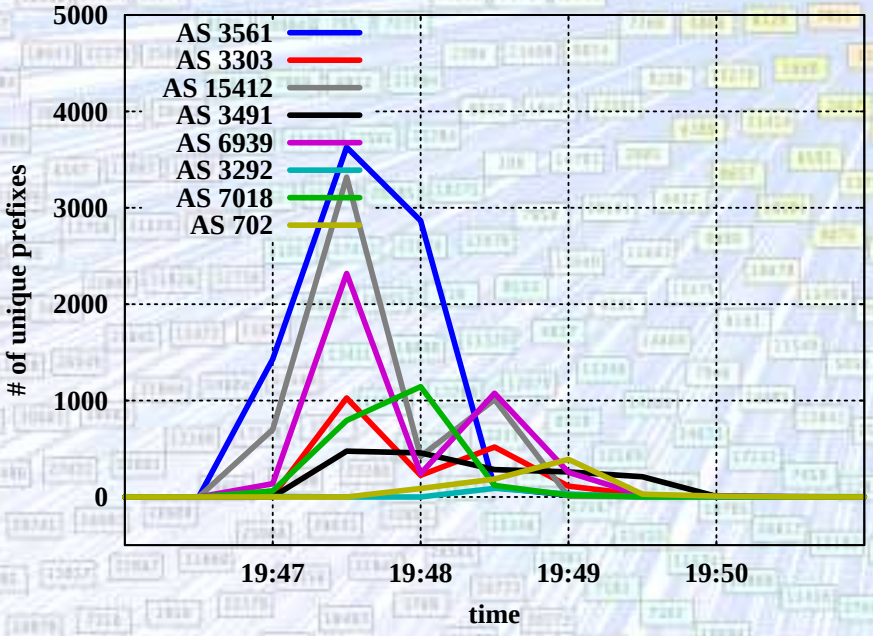

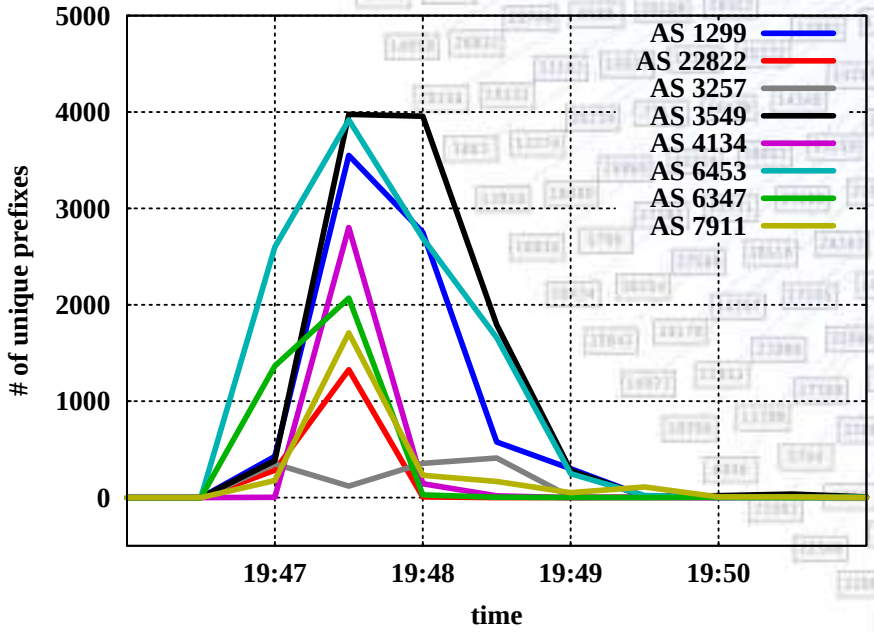

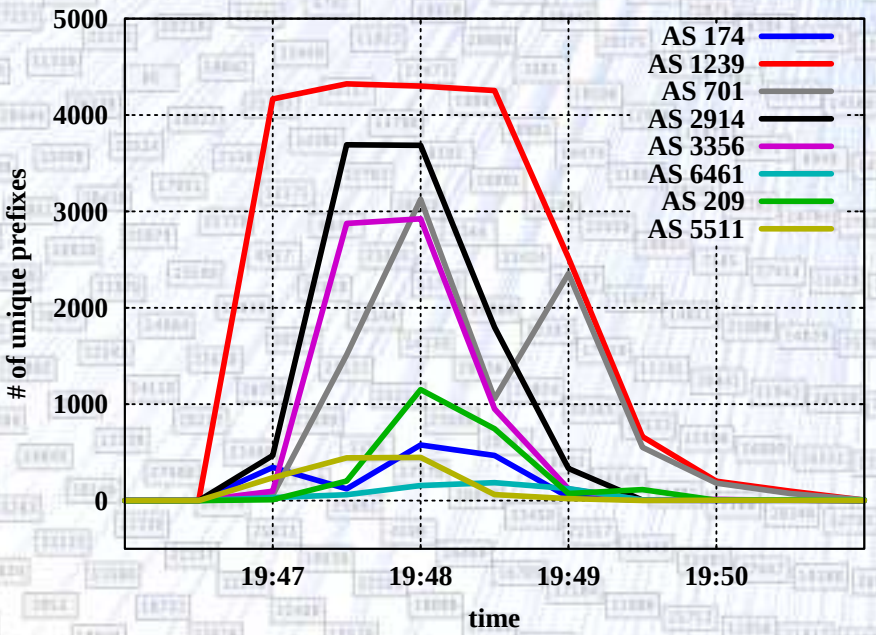

Rates of Advertisement – Event #2

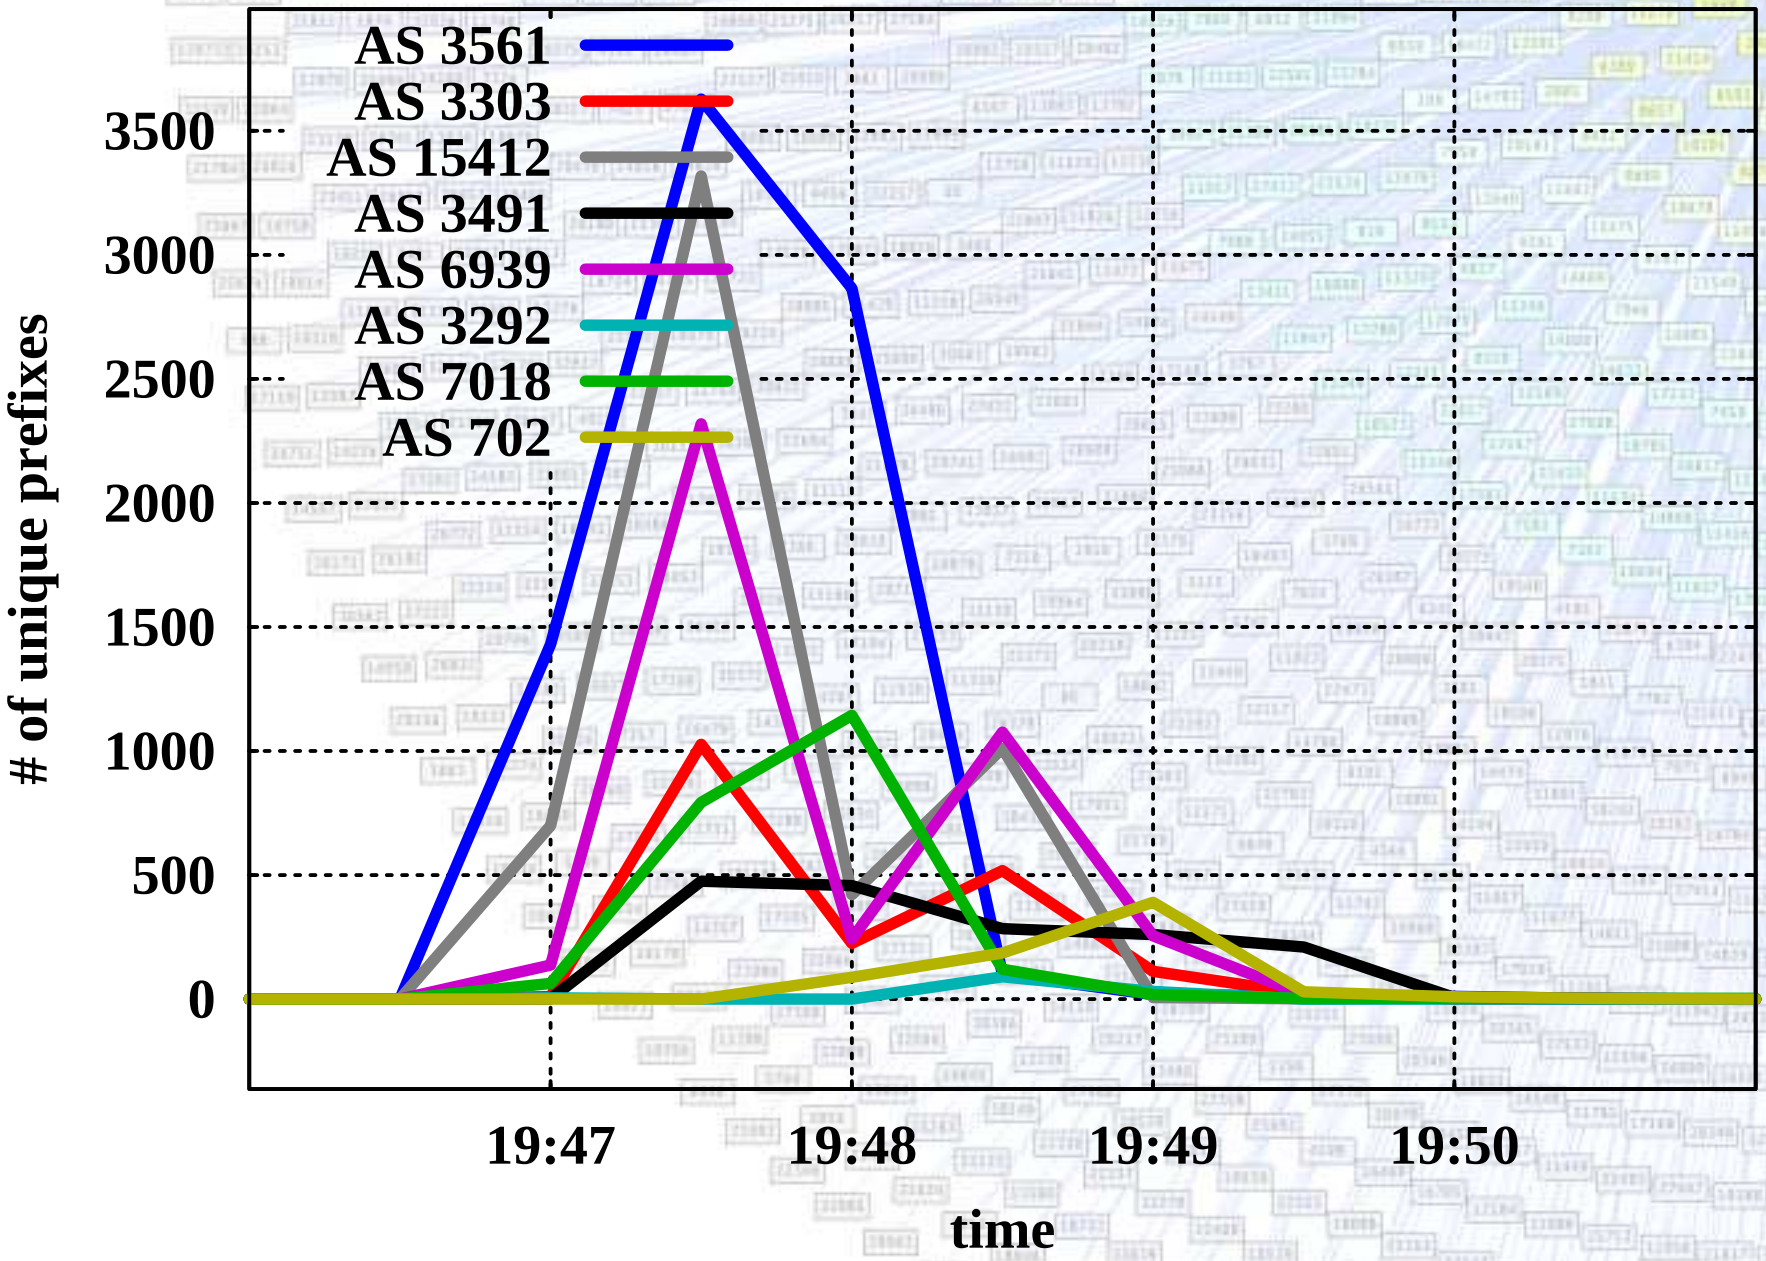

Rates of Advertisement – Event #2

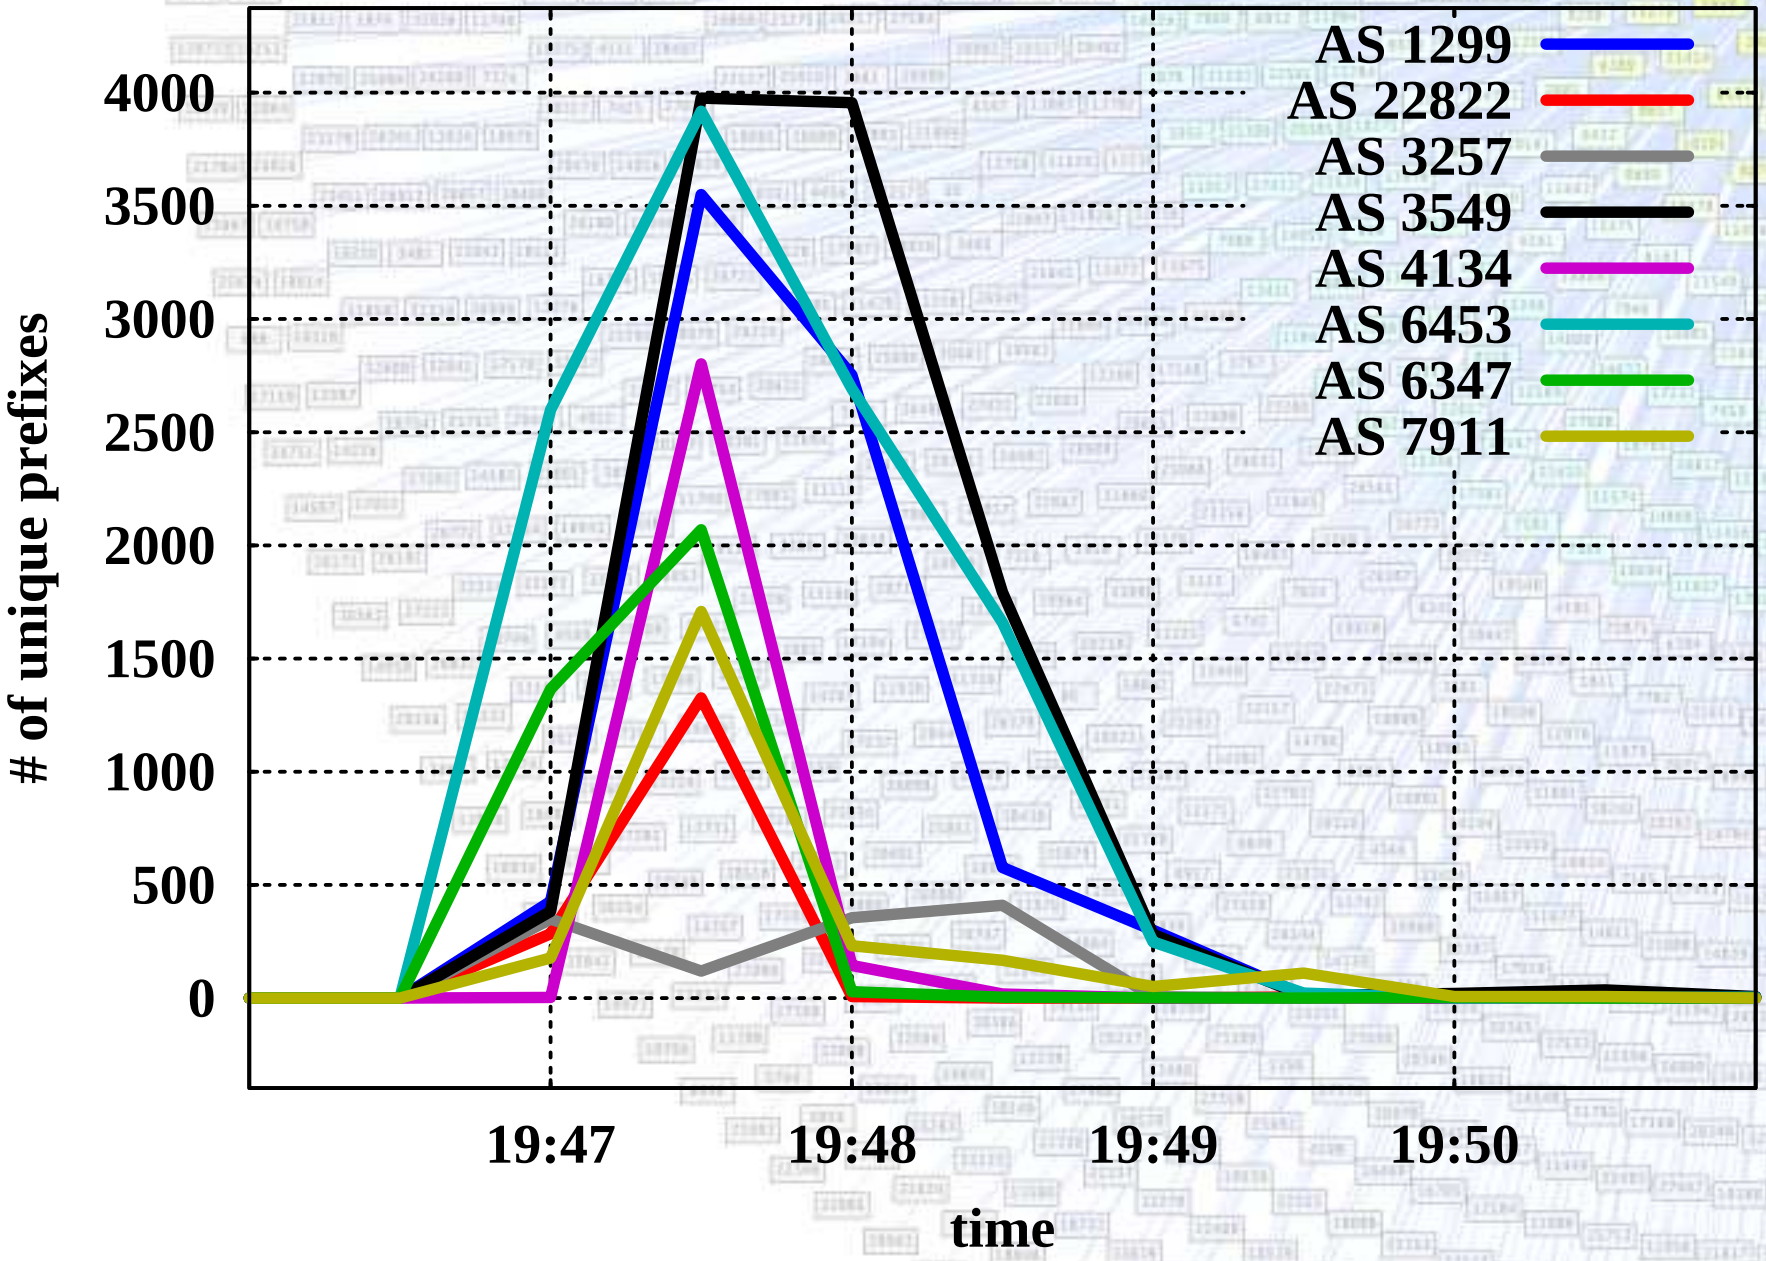

Rates of Advertisement – Event #2

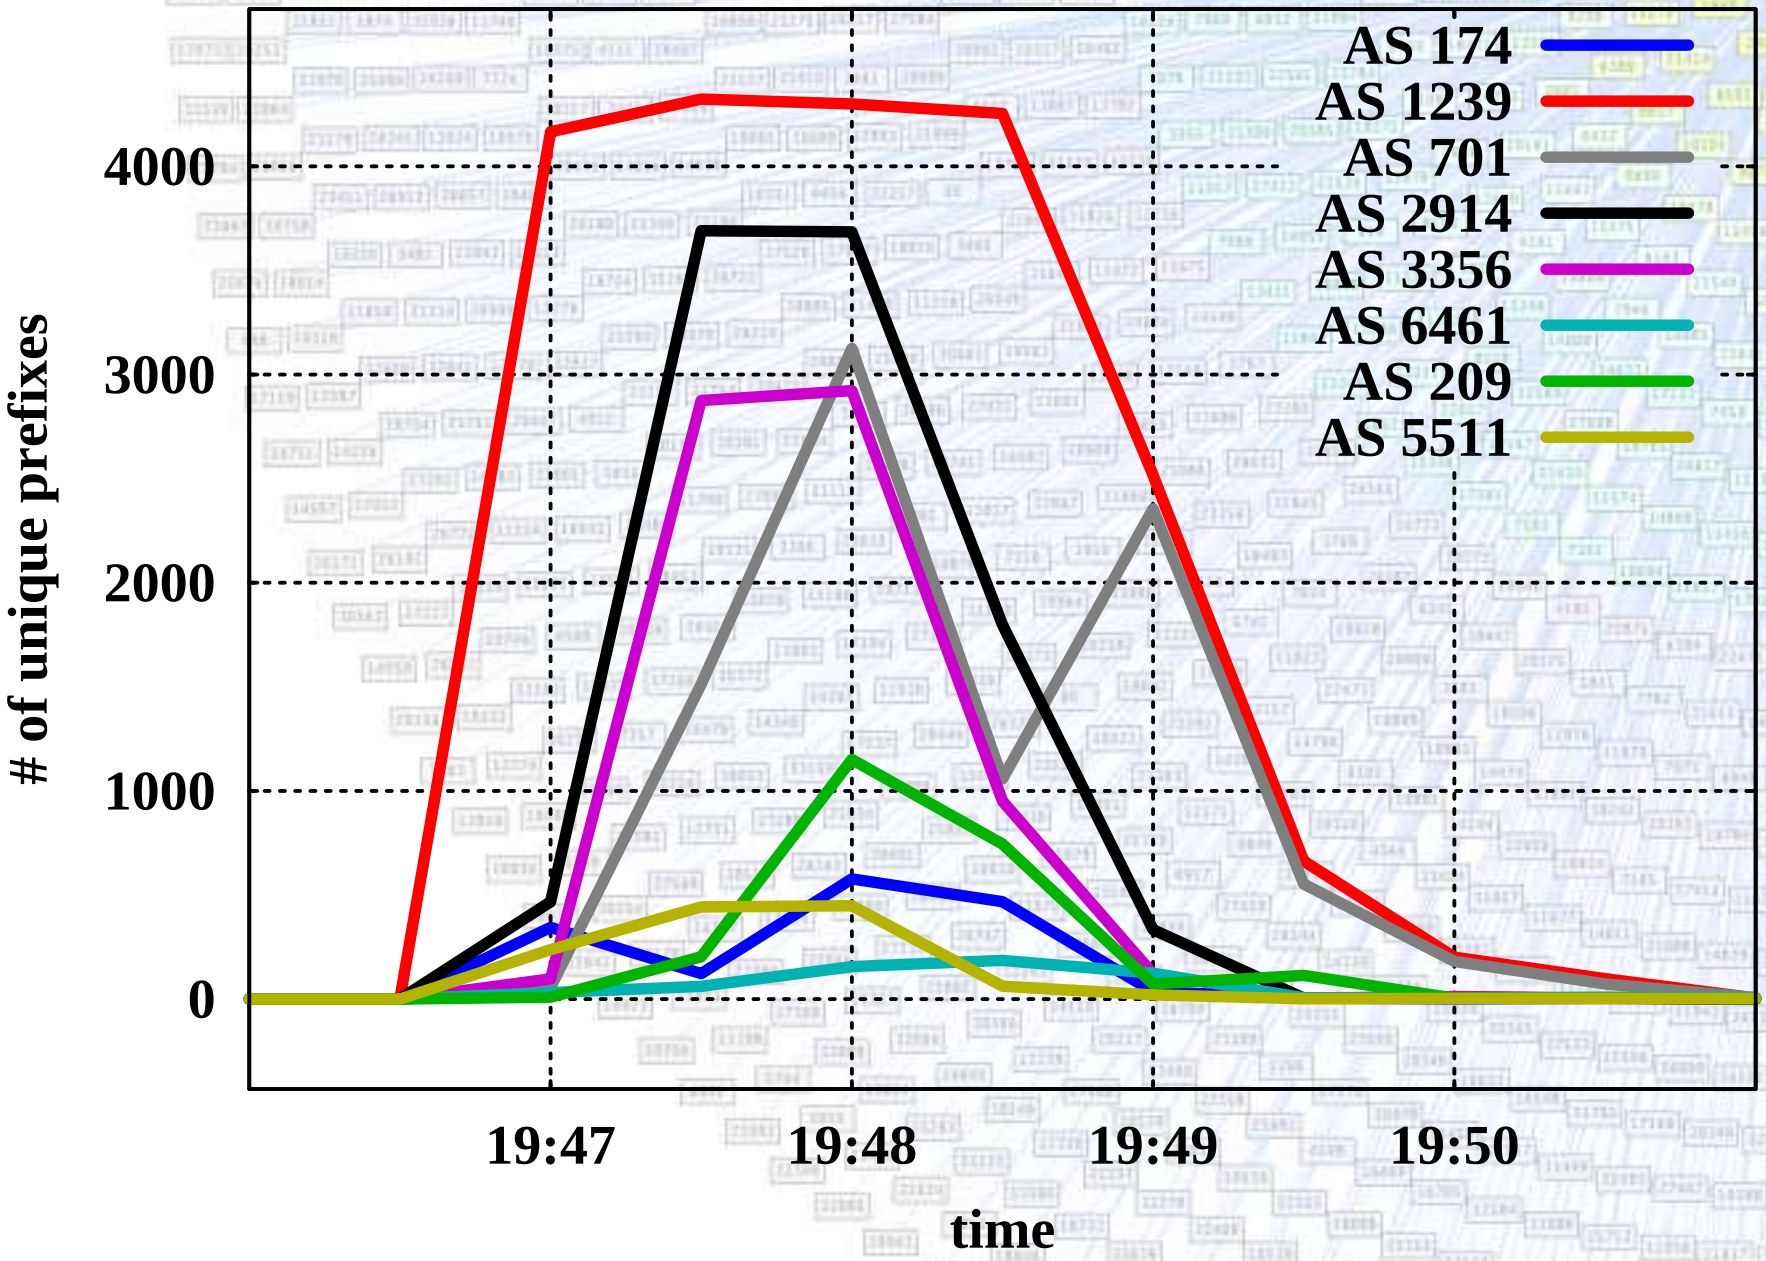

Prefixes Carried – Event #2

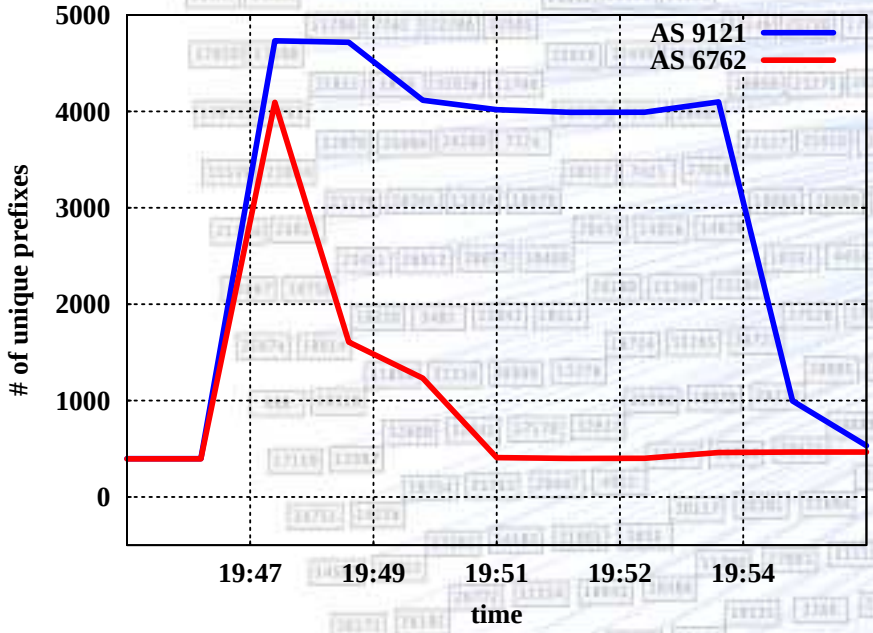

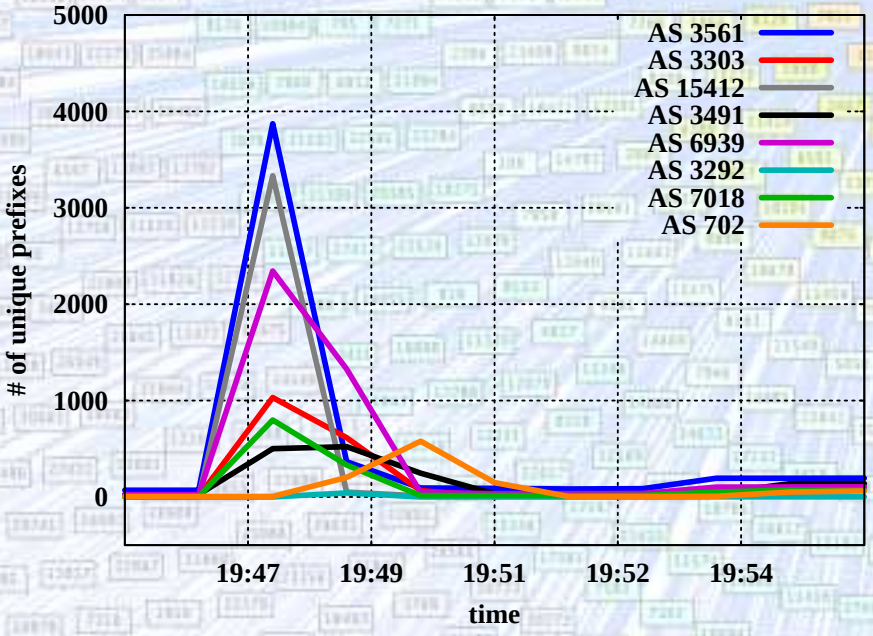

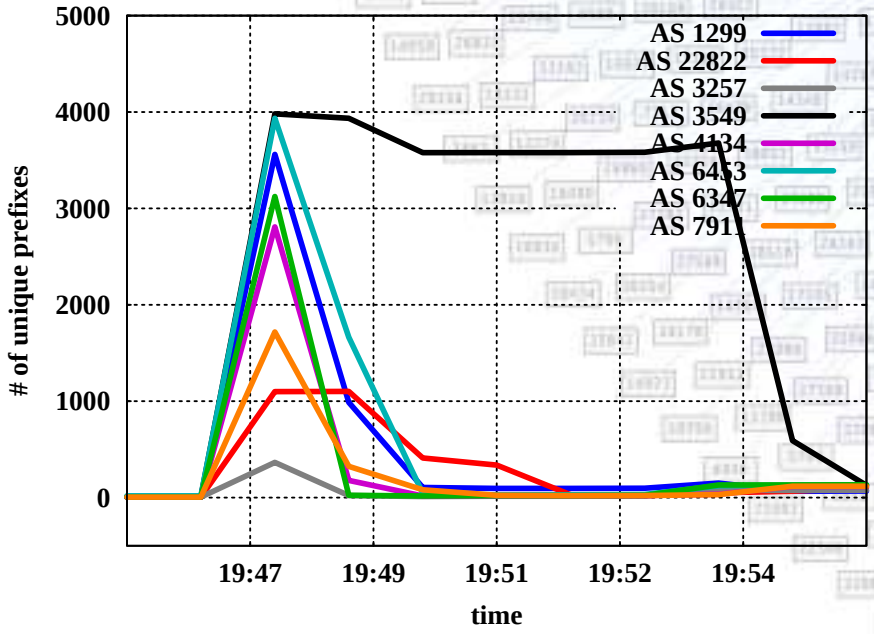

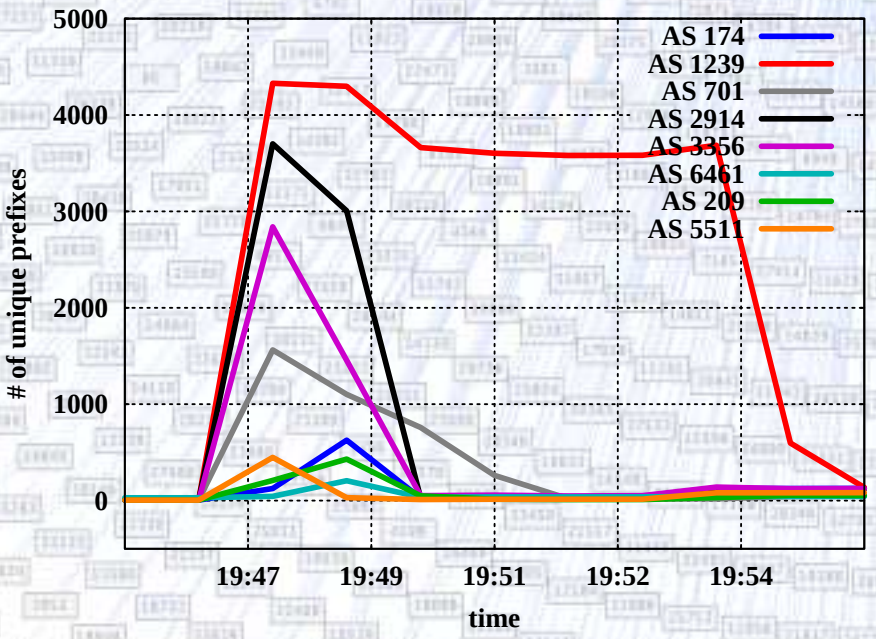

Prefixes Carried – Event #2

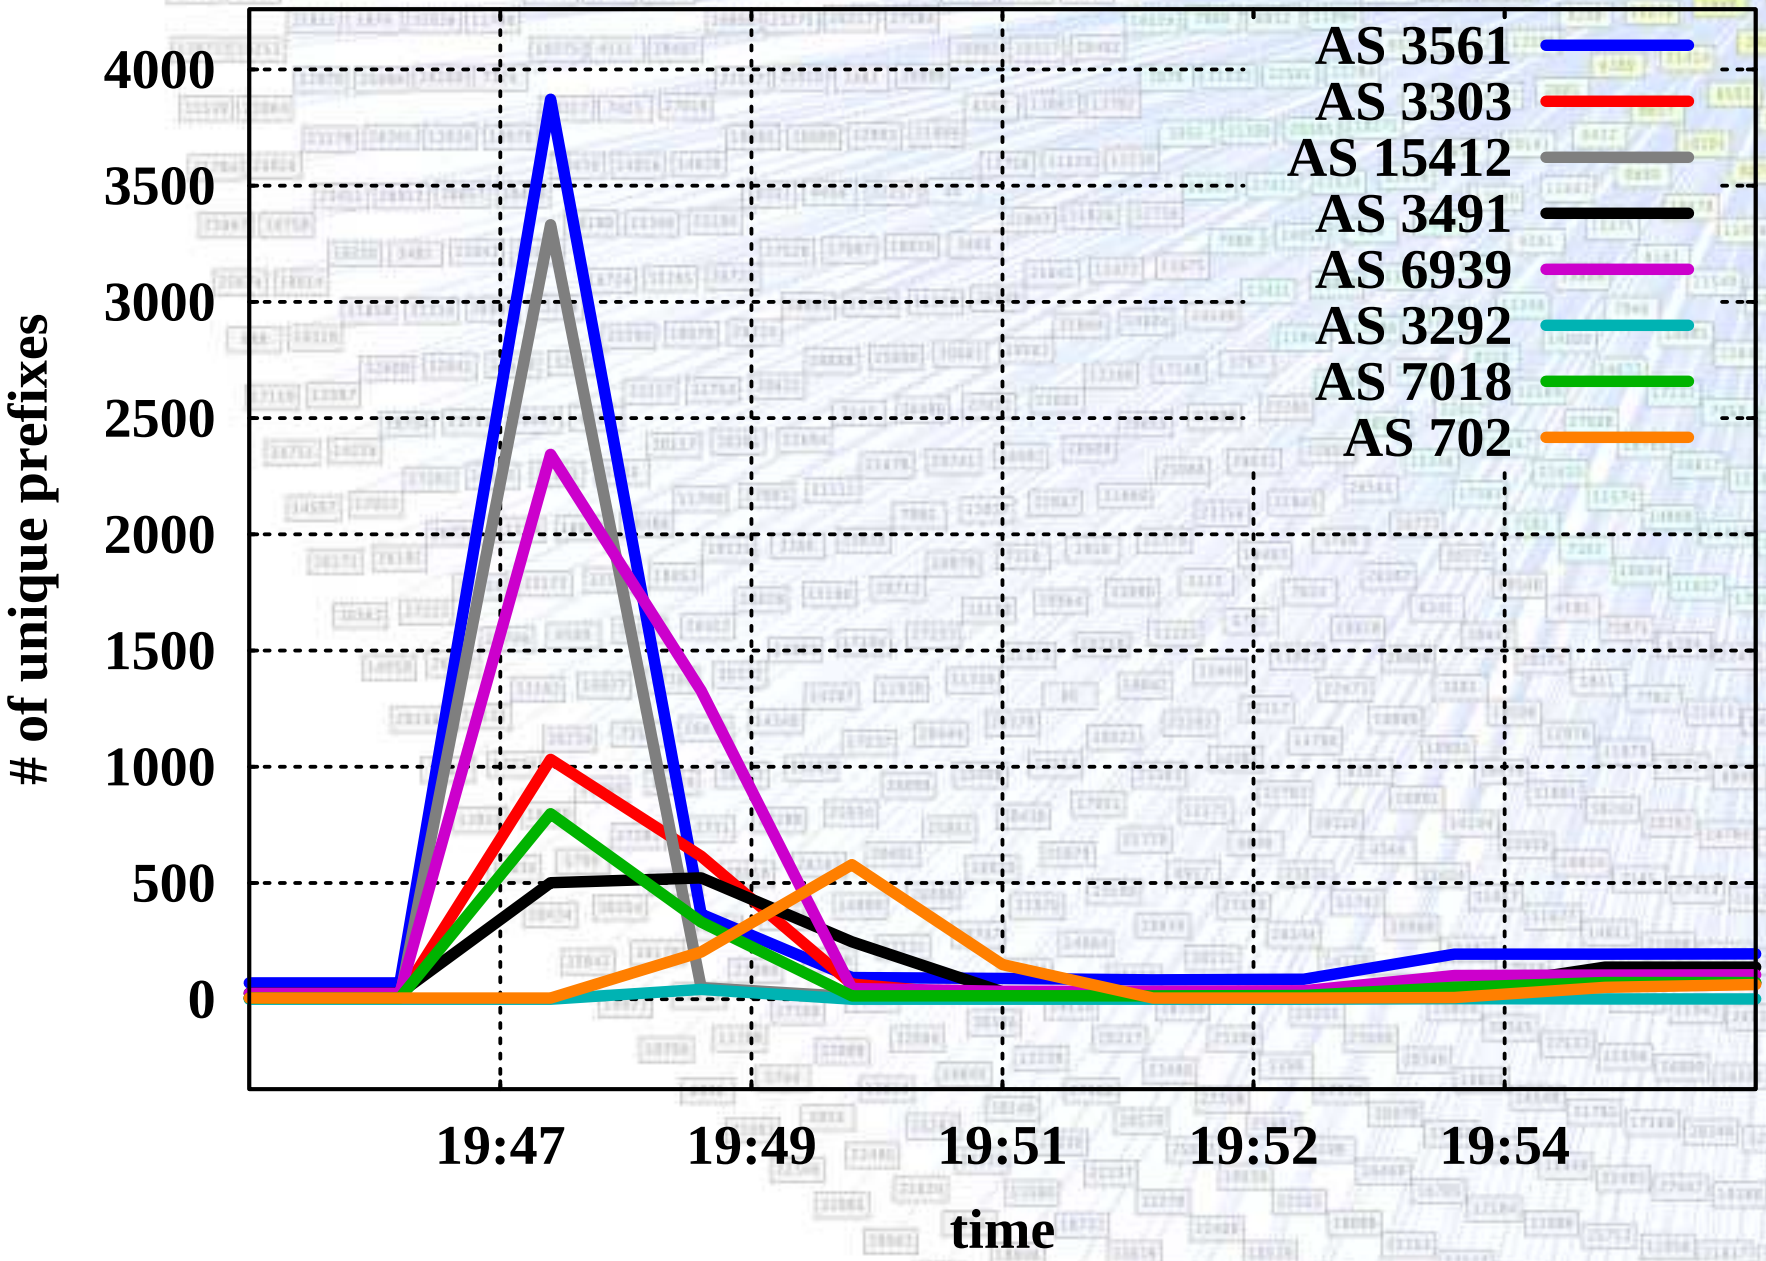

Prefixes Carried – Event #2

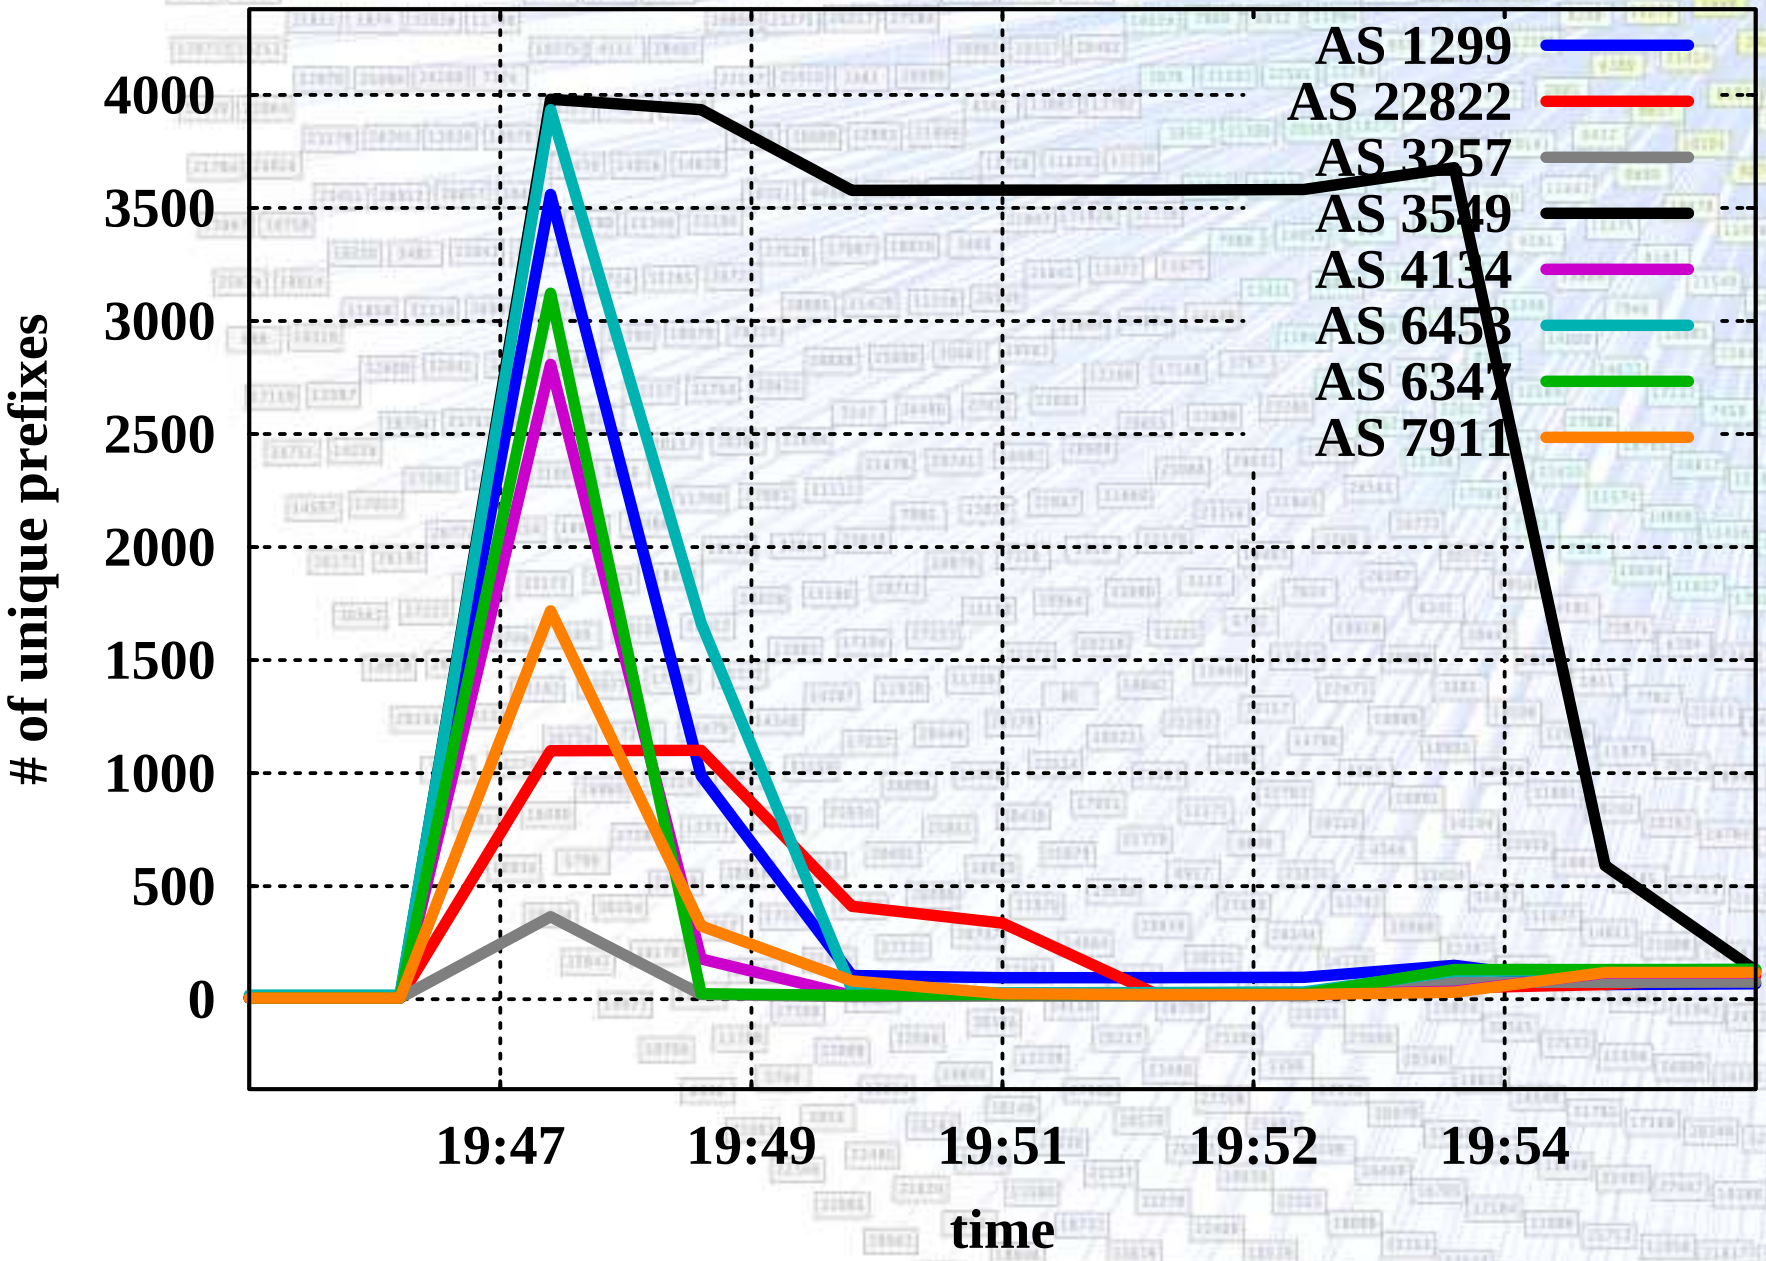

Prefixes Carried – Event #2